

Parameterized Pattern Matching – Succinctly

Arnab Ganguly*

Rahul Shah†

Sharma V. Thankachan‡

April 7, 2016

Abstract

The fields of succinct data structures and compressed text indexing have seen quite a bit of progress over the last 15 years. An important achievement, primarily using techniques based on the Burrows-Wheeler Transform (BWT), was obtaining the full functionality of the suffix tree in the optimal number of bits. A crucial property that allows the use of BWT for designing compressed indexes is *order-preserving suffix links*. Specifically, the relative order between two suffixes in the subtree of an internal node is same as that of the suffixes obtained by truncating the first character of the two suffixes. Unfortunately, in many variants of the text-indexing problem, for e.g., parameterized pattern matching and 2D pattern matching, this property does not hold. Consequently, the compressed indexes based on BWT do not directly apply. Furthermore, a compressed index for any of these variants has been elusive throughout the advancement of the field of succinct data structures. We achieve a positive breakthrough on one such problem, namely the *Parameterized Pattern Matching* problem.

Let T be a text that contains n characters from an alphabet Σ , which is the union of two disjoint sets: Σ_s containing static characters (s-characters) and Σ_p containing parameterized characters (p-characters). A pattern P (also over Σ) matches an equal-length substring S of T iff the s-characters match exactly, and there exists a one-to-one function that renames the p-characters in S to that in P . The task is to find the starting positions (occurrences) of all such substrings S . Previous indexing solution [Baker, STOC 1993], known as *Parameterized Suffix Tree*, requires $\Theta(n \log n)$ bits of space, and can find all *occ* occurrences in time $\mathcal{O}(|P| \log \sigma + \text{occ})$, where $\sigma = |\Sigma|$. In this paper, we present the first succinct index for this problem. Specifically, we introduce two indexes with the following space-and-time tradeoffs:

- $\mathcal{O}(n \log \sigma)$ bits and $\mathcal{O}(|P| \log \sigma + \text{occ} \cdot \log n)$ query time.
- $n \log \sigma + \mathcal{O}(n)$ bits and $\mathcal{O}((|P| + \text{occ} \cdot \log n) \log \sigma \log \log \sigma)$ query time.

Furthermore, the techniques are extended to obtain the first succinct representation of the index of Shibuya for *Structural Matching* [SWAT, 2000], and of Idury and Schäffer for *Parameterized Dictionary Matching* [CPM, 1994].

Keywords: Succinct data structures, Suffix trees, Burrows-Wheeler Transform

*School of Electrical Engineering and Computer Science. Louisiana State University, Baton Rouge, USA. **email:** agangu4@lsu.edu

†School of Electrical Engineering and Computer Science. Louisiana State University, Baton Rouge, USA. Also affiliated to: National Science Foundation, USA. **email:** rahul@csc.lsu.edu, rahul@nsf.gov

‡School of Computational Science and Engineering. Georgia Institute of Technology, Atlanta, USA. **email:** sharma.thankachan@gatech.edu

1 Introduction

Pattern Matching is the algorithmic framework of finding all starting positions (or simply, occurrences) of a pattern in a text. In earlier works [BM77, KR87, KJP77], both text and pattern were provided as input. In most cases, however, the text does not change and patterns come in as online query. This motivated the development of full text indexes i.e., pre-process the text and then build a data structure. Using the data structure, whenever a pattern comes as query, we can efficiently report all occurrences without having to read the entire text. Here and henceforth, T is a text containing n characters chosen from an alphabet Σ of size $\sigma \geq 2$, P is a pattern containing $|P|$ characters (also from Σ), and occ is the number of occurrences of P in T . *Suffix Trees* [Ukk95, Wei73] and *Suffix Arrays* along with *Longest Common Prefix* (LCP) array [MM93] can answer pattern matching queries in $\mathcal{O}(|P| + occ)$ and in $\mathcal{O}(|P| + \log n + occ)$ time respectively. Both suffix trees and suffix arrays require $\Theta(n \log n)$ bits of space, which is too large for most practical purposes (15-50 times the text). Grossi and Vitter [GV00], and Ferragina and Manzini [FM00] were the first to address this problem. Their respective space efficient indexes, namely *Compressed Suffix Arrays* (CSA) and *FM-Index*, have found a lot of applications, and led to the establishment of the field of compressed text indexing. See [NM07] for a comprehensive survey.

Suffix trees have myriad applications (see Dan Gusfield’s book [Gus97]). For many applications, suffix trees have to be augmented with additional data. After the initial advent of compressed text indexing structures, *compressed suffix trees* [RNO11, Sad07] replaced the older suffix tree as a black-box. The focus was now on how to compress the augmenting data in informational theoretic minimal space. Probably the first such work was that by Sadakane [Sad02b] who showed that given a (compressed) suffix array, the LCP array can be maintained in an additional $2n + o(n)$ bits. Fischer [Fis10] improved this to $o(n)$ bits. The *Range Minimum Query* (RMQ) data structure is an integral component to many applications related to suffix trees/arrays. Fischer and Heun [FH07] showed how to maintain the RMQ data structure in $2n + o(n)$ bits. In another direction, Sadakane [Sad02a] considered the problem of document retrieval, and presented succinct indexes. Later, other improvements (see [Nav13] for a survey) culminated in the best known document listing index of Tsur [Tsu13]. Other problems such as property matching [HPST13] and matching statistics [OGK10] also fall under the same theme.

To the best of our knowledge, no progress has been achieved on designing space efficient indexes (analogous to that of the Compressed Suffix Array and FM-index) for addressing the second kind of applications [Bak93, Gia93, KEF⁺14, Shi00] which require variants of the classical suffix tree. We now formally define such a problem, that of parameterized pattern matching.

Problem 1 ([Bak93]) *Let Σ be an alphabet of size $\sigma \geq 2$, which is the union of two disjoint sets: Σ_s having σ_s static characters (s -characters) and Σ_p having σ_p parameterized characters (p -characters). Two strings, having equal length, are a parameterized match (p -match) if one can be transformed to the other by applying a one-to-one function that renames the characters in Σ_p . Let T be a text having n characters chosen from Σ . We assume that T terminates in an s -character $\$$ that appears only once. The task is to index T , such that for a pattern P (also over Σ), we can report all starting positions (occurrences) of the substrings of T that are a p -match with P .*

Consider the following example. Let $\Sigma_s = \{A, B, C, \$\}$ and $\Sigma_p = \{w, x, y, z\}$. Then, $P = AxByCx$ p -matches within $T[1, 21] = AyBxCyAwBxCzxyAzBwCz\$$ at positions 1 and 15. At position 1,

the mapping is $x \rightarrow y$ and $y \rightarrow x$; whereas at position 15, the mapping is $x \rightarrow z$ and $y \rightarrow w$. Note that P does not match at position 7 because x would have to match with both w and z .

1.1 Applications

Baker [Bak93] introduced this problem, with the main motivation being software plagiarism detection. They introduced a $\Theta(n \log n)$ -bit index, known as the *Parameterized Suffix Tree* (p-suffix tree), that can report all *occ* p-matches in $\mathcal{O}(|P| \log \sigma + \text{occ})$ time. The role of the problem and the efficiency of p-suffix trees was presented using a program called *Dup* [Bak95]. Subsequently, the methodology became an integral part of various tools for software version management and clone detection, where identifiers and/or literals are renamed. The most state-of-the-art approaches [BYdM⁺98, KFF06, RD04] to clone detection use a hybrid approach, such as a combination of (i) a parse tree, which converts literals into parameterized symbols, and (ii) a p-suffix tree on top of these symbols. Unfortunately, as with traditional suffix trees, the space occupied by p-suffix trees is too large for most practical purposes, bringing in the demand for space-efficient variants. In fact, one of the available tools (CLICS [CLI]) very clearly acknowledges that the major space consumption is due to the use of p-suffix tree. Some other tools [HJHC10] use more IR type methodology for indexing repositories based on variants of the inverted index. Although less space consuming, there are no theoretical guarantees possible on query-times in such indexes. Since p-suffix trees accommodate substring match, tools based on them can also be powerful in recognizing smaller code fragment clones. Following are a few other works that have used p-suffix trees: finding relevant information based on regular expressions in sequence databases [dMRS05, dMRS07], detecting cloned web pages [DLDPFG01], detecting similarities in JAVA sources from bytecodes [BM98], etc. Further generalizations of p-matching have also played an important role in computational biology for finding similar sequences [AAC⁺03, Shi00].

1.2 Contribution

An important ingredient of suffix trees, crucial to compressed text indexing, is *suffix links*. Specifically, let u be a non-root node in the suffix tree, and let the string corresponding to the path from root to u be S . Then, the node whose path label is obtained by truncating the first character of S necessarily exists. Let v be that node. The suffix link of u points to v . In suffix trees, the leaves are arranged in the lexicographic order of the suffix they represent. Suffix links have the following so called *order-preserving* property. The leaves obtained by following suffix links from the leaves in u 's subtree appear in the same relative order in the subtree of v . Thus, the permutation of suffixes in v 's subtree can be encoded in terms of the permutation in u 's subtree. In applications like p-suffix tree [Bak93], 2D suffix tree [Gia93, KKP03], structural suffix tree [Shi00] etc., this property is not true. This brings in new challenges in how to encode such permutations, and whether succinct data structures are even possible. By presenting the first succinct index for Problem 1, we affirmatively answer this in the case of p-pattern matching. The following theorem summarizes our main contribution.

Theorem 1 *All occ positions where P is a p-match with T can be found as follows:*

- (a) *in $\mathcal{O}(|P| \log \sigma + \text{occ} \cdot \log n)$ time using an $\mathcal{O}(n \log \sigma)$ -bit index.*
- (b) *in $\mathcal{O}((|P| + \text{occ} \cdot \log n) \log \sigma \log \log \sigma)$ time using an $n \log \sigma + \mathcal{O}(n)$ -bit index.*

We introduce a BWT-like transform for a parameterized text, called the *Parameterized BWT*. Using this, we show how to handle the *order-inversion* in the case of p-suffixes when their first characters are truncated. To achieve this, we implement analogous versions of the last-to-first column mapping and backward-search technique of Ferragina and Manzini [FM00]. This encompasses the major contribution of our work, which is achieved using newly introduced concepts, and existing succinct data structure toolkit. The techniques are then extended to address two related problems. Specifically, we show that an index with the same space-and-time tradeoff (as in Theorem 1) can be obtained to simulate the functionalities of the structural suffix tree (a generalization of the p-suffix tree) of Shibuya [Shi00]. A succinct representation of the parameterized dictionary matching index of Idury and Schäffer [IS94] is also achieved.

1.3 Previous and Related Work

Parameterized pattern matching has seen constant development since its inception by Baker [Bak93] in 1993. The $\mathcal{O}(n\sigma_p + n \log \sigma)$ time algorithm presented in [Bak93] for constructing p-suffix trees was improved to $\mathcal{O}(n \log \sigma)$ by Kosaraju [Kos95]. Later, Cole and Hariharan [CH03] gave an $\mathcal{O}(n)$ time randomized algorithm. Amir, Farach, and Muthukrishnan [AFM94] presented an algorithm when the text T and pattern P are both provided at query time. More recently, Jalsenius, Porat and Sach presented a solution for p-matching in the streaming model [JPS13]. Other variants include the two-dimensional p-matching problem [AAC⁺03], approximate p-matching [HLS04], parameterized longest common subsequence [KKL09], and p-matching on non-linear structures [AN09]. We refer the reader to [Lew15, MP15] for recent surveys.

1.4 Map

In Section 2, we first revisit a few standard text indexing data structures; in particular, we take a close look at the parameterized suffix tree of Baker [Bak93] as it plays a crucial role in the proposed index. In Section 3, we introduce the parameterized BWT. In Section 4, we implement an analogous version of the last-to-first column mapping (LF mapping), introduced in the FM-index of Ferragina and Manzini [FM00]. We also present the key components of the proposed index in this section. Using the LF mapping implementation and its accompanying data structures, an adaptation of the backward search technique of Ferragina and Manzini [FM00] enables us to arrive at Theorem 1a. Section 5 presents the details. We prove Theorem 1b in Section 6. Section 7 is devoted to obtain a succinct representation of the structural suffix tree of Shibuya [Shi00]. Section 8 presents a succinct representation of the parameterized dictionary matching index of Idury and Schäffer [IS94]. Finally, by presenting a few unaddressed problems, we conclude the paper in Section 9.

2 Background

Throughout this paper, we use the following terminologies: for a string S , $|S|$ is its length, $S[i]$, $1 \leq i \leq |S|$, is its i th character and $S[i, j]$ is its substring from i to j (both inclusive). If $i > j$, $S[i, j]$ denotes an empty string. Also S_i denotes the circular suffix starting at position i . Specifically, S_i is S if $i = 1$ and is $S[i, |S|] \circ S[1, i - 1]$ otherwise, where \circ denotes the *concatenation*.

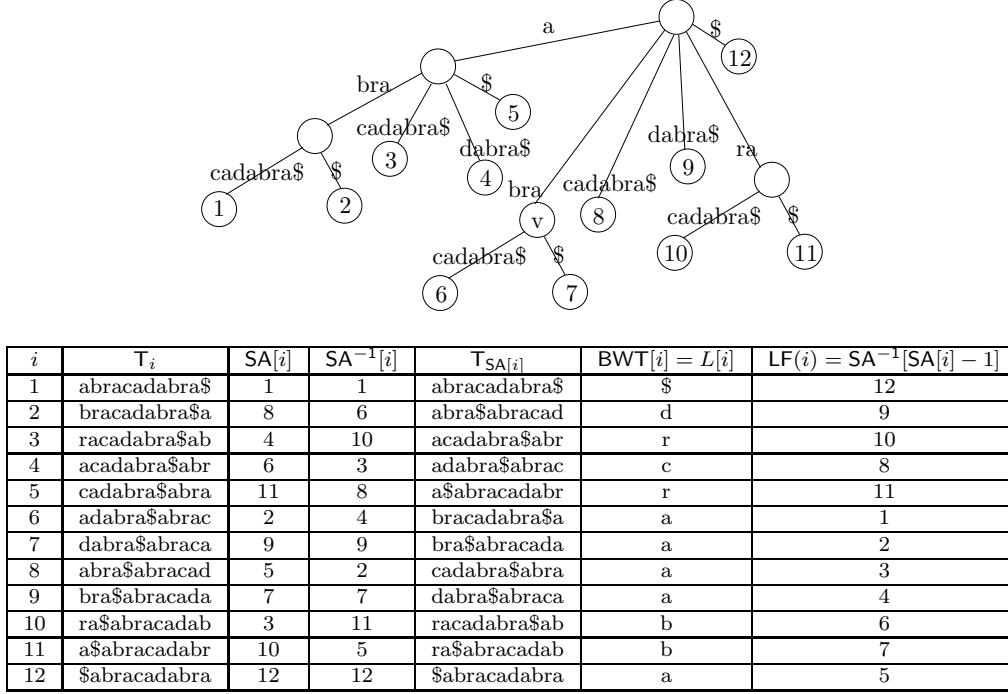


Figure 1: The text is $T[1, 12] = abracadabra\$$. We assume the lexicographic order is $a \prec b \prec c \prec d \prec r \prec \$$. Here, locus of br is the node v and suffix range of br is $[6, 7]$.

2.1 Suffix Tree and Suffix Array

A suffix tree ST is a compact trie that stores all the (non-empty) suffixes of T [Ukk95, Wei73]. Leaves in the suffix tree are numbered in the lexicographic order of the suffix they represent. The locus of a pattern P is the highest node u such that P is a prefix of the string formed by concatenating the edge labels from the root to u . The suffix range of P is the set of leaves in the subtree of ST rooted at the locus of P . See Figure 1 for an illustration. Using suffix trees, under the perfect hashing assumption, the locus node (or equivalently, the suffix range) of any pattern P can be computed in time $O(|P|)$. Without hashing, the suffix range can be computed in time $O(|P| \log \sigma)$ by maintaining the outgoing edges in sorted order of the first character on the edge.

A suffix array SA is an array of length n that maintains the lexicographic arrangement of all the suffixes of T [MM93]. More specifically, if the i th smallest suffix of T starts at j , then $SA[i] = j$ and $SA^{-1}[j] = i$. The former is referred to as the *suffix array value* and the latter as the *inverse suffix array value*. See Figure 1 for an illustration. Once the suffix range $[sp, ep]$ of P is found, by using the suffix array, we can find the set $\{SA[i] \mid sp \leq i \leq ep\}$ of all *occ* occurrences of P . The time needed is $O(|P| + occ)$.

2.2 Burrows-Wheeler Transform and FM-Index

Burrows and Wheeler [BW94] introduced a reversible transformation of the text, known as the Burrows-Wheeler Transform (BWT). Recall that T_x is the circular suffix starting at position x . Then, BWT of a text T is obtained as follows: first create a conceptual matrix M , such that each

row of M corresponds to a unique circular suffix, and then lexicographically sort all rows. The BWT of the text T is the last column L of M . See Figure 1 for an illustration. Note that BWT is essentially a permutation of the text T .

2.2.1 Last-to-First Column Mapping

The underlying principle that enables pattern matching using an FM-Index [FM00] is the last-to-first column mapping. For any $i \in [1, n]$, $\text{LF}(i)$ is the row j in the matrix M where $\text{BWT}[i]$ appears as the first character in $T_{\text{SA}[j]}$. Specifically, $\text{LF}(i) = \text{SA}^{-1}[\text{SA}[i] - 1]$. Once the BWT of the text is obtained, $\text{LF}(i)$ for any suffix i is computed by the following formula:

$$\text{LF}(i) = \text{Count}_{<}[\text{BWT}[i]] + \text{Count}_{=}[\text{BWT}[i], i]$$

where $\text{Count}_{<}[c]$ is the total number of characters in T smaller than c , and $\text{Count}_{=}[c, j]$ is the number of occurrences of c in $\text{BWT}[1, j]$. By maintaining a Wavelet Tree [GGV03] over $\text{BWT}[1, n]$, both $\text{Count}_{<}[\cdot]$ and $\text{Count}_{=}[c, \cdot]$ operations can be carried out in $\mathcal{O}(\log \sigma)$ time. See Section 4.1.1 for more details on the Wavelet Tree data structure.

2.2.2 Backward Search

Ferragina and Manzini [FM00] showed that using LF-mapping, the suffix range $[\text{sp}, \text{ep}]$ of a pattern $P[1, |p|]$ can be found by reading P starting from the last character. Specifically, for $i > 1$, suppose the suffix range of $P[i, p]$ is known. Then, the suffix range of $P[i - 1, p]$ can be obtained using LF-mapping. The backward search procedure is summarized in Algorithm 1.

Algorithm 1 for finding Suffix Range of $P[1, p]$

```

1:  $c \leftarrow P[p], i \leftarrow p$ 
2:  $\text{sp} \leftarrow 1 + \text{Count}_{<}[c], \text{ep} \leftarrow \text{Count}_{<}[c + 1]$ 
3: while ( $\text{sp} \leq \text{ep}$  and  $i \geq 2$ ) do
4:    $c \leftarrow P[i - 1]$ 
5:    $\text{sp} \leftarrow 1 + \text{Count}_{<}[c] + \text{Count}_{=}[c, \text{sp} - 1]$ 
6:    $\text{ep} \leftarrow \text{Count}_{<}[c] + \text{Count}_{=}[c, \text{ep}]$ 
7: end while
8: if ( $\text{sp} < \text{ep}$ ) then “no match found” else return  $[\text{sp}, \text{ep}]$ 
```

By maintaining a sampled-suffix array, which occupies $o(n)$ bits in total, each occurrence can be reported in $\mathcal{O}(\log^{1+\epsilon} n \log \sigma)$ time, where $\epsilon > 0$ is an arbitrarily small constant. The key idea is to explicitly store SA for some specific values. Then, a particular $\text{SA}[i]$ can be obtained directly if the value has been explicitly stored; otherwise, it can be obtained via at most $\lceil \log^{1+\epsilon} n \rceil$ successive $\text{LF}(\cdot)$ operations.

2.3 Baker’s Parameterized Suffix Tree

Baker [Bak93] introduced the following encoding scheme for matching strings over a parameterized alphabet $\Sigma = \Sigma_s \cup \Sigma_p$. A string S is encoded into a string $\text{prev}(S)$ of length $|S|$ by replacing the first occurrence of every p-character in S by 0 and any other occurrence of a p-character by the difference

in text position from its previous occurrence. More specifically, for any $i \in [1, |S|]$, $\text{prev}(S)[i] = S[i]$ if $S[i]$ is an s-character; otherwise, $\text{prev}(S)[i] = (i - j)$, where $j < i$ is the last occurrence (if any) of $S[i]$ before i . If j does not exist, then $j = i$. For example, $\text{prev}(AxB y Cx) = A0B0C4$, where $A, B \in \Sigma_s$ and $x, y \in \Sigma_p$. Note that $\text{prev}(S)$ is a string over $\Sigma' = \Sigma_s \cup \{0, 1, \dots, |S| - 1\}$, and can be computed in time $\mathcal{O}(|S| \log \sigma)^1$.

Fact 1 ([Bak93]) *Two strings S and S' are a p -match iff $\text{prev}(S) = \text{prev}(S')$. Also S and a prefix of S' are a p -match iff $\text{prev}(S)$ is a prefix of $\text{prev}(S')$.*

Moving forward, we follow the convention below.

Convention 1 *In Σ' , the integer characters (corresponding to p -characters) are lexicographically smaller than s -characters². An integer character i comes before another integer character j iff $i < j$. Also, $\$$ is lexicographically larger than all other characters.*

Parameterized Suffix Tree (pST) is the compacted trie of all strings in $\mathcal{P} = \{\text{prev}(T[k, n]) \mid 1 \leq k \leq n\}$. Each edge is labeled with a string over Σ' . We use $\text{path}(u)$ to denote the concatenation of edge labels on the path from root to node u . Clearly, pST consists of n leaves (one per each encoded suffix) and at most $n - 1$ internal nodes. The space required is $\Theta(n \log n)$ bits. See Figure 2 for an illustration. The path of each leaf node corresponds to the encoding of a unique suffix of T , and leaves are ordered in the lexicographic order of the corresponding encoded suffix.

To find all occurrences of P , traverse pST from root by following the edges labels and find the highest node u (called *locus*) such that $\text{path}(u)$ is prefixed by $\text{prev}(P)$. Then find the range $[\text{sp}, \text{ep}]$ (called *suffix range* of $\text{prev}(P)$) of leaves in the subtree of u and report $\{\text{pSA}[i] \mid \text{sp} \leq i \leq \text{ep}\}$ as the output. Here, $\text{pSA}[1, n]$ is the parameterized suffix array i.e., $\text{pSA}[i] = j$ and $\text{pSA}^{-1}[j] = i$ iff $\text{prev}(T[j, n])$ is the i th lexicographically smallest string in \mathcal{P} . Note that $\text{path}(\ell_i) = \text{prev}(T[\text{pSA}[i], n])$, where ℓ_i is the i th leftmost leaf in pST. The query time is $\mathcal{O}(|P| \log \sigma + \text{occ})^3$.


3 Parameterized Burrows-Wheeler Transform

We introduce a similar transform to that of the BWT, which we call the Parameterized Burrows-Wheeler Transform (p-BWT). To obtain the p-BWT of T , we create a similar matrix M as in the case of the BWT described earlier. Then, we sort all this rows lexicographically according to the $\text{prev}(\cdot)$ encoding of the corresponding unique circular suffix, and obtain the last column L of the

¹ Assume that Σ comprises of integers from 1 through σ_p representing the p -characters and $\sigma_p + 1$ through σ representing the s -characters. Maintain an array $F[1, \sigma_p]$, all values initialized to 0. Scan S from left to right, and suppose a p -character $k \leq \sigma_p$ appears at position x . Then $\text{prev}(S)[x]$ equals 0 if $F[k] = 0$, and equals $x - F[k]$, otherwise. Update $F[k] = x$. Time is $\mathcal{O}(|S|)$. For a general Σ , we take the convention that all p -characters are lexicographically smaller than s -characters. Then, the string S over Σ can be converted into a string over an integer alphabet (as described previously) using a balanced binary search tree in $\mathcal{O}(|S| \log \sigma)$ time by simply mapping the k th smallest symbol to the integer k .

² If Σ_s contains integers, then we simply map the k th smallest s -character to the integer $|S| + k$ to ensure that it is disjoint from the set $\{0, 1, \dots, |S| - 1\}$. The string $|S|$ can be pre-processed in $\mathcal{O}(|S| \log \sigma)$ time using a balanced binary search tree to ensure that this condition holds. This pre-processing will not affect any time complexities.

³ The query time is achieved by using perfect hashing, i.e., at each node, we can navigate to its appropriate child in constant time. The $\mathcal{O}(|P| \log \sigma)$ factor is due to the time to obtain $\text{prev}(P)$. If we remove perfect hashing, then the time increases to $\mathcal{O}(|P| \log n + \text{occ})$, as each internal node may have $\Theta(n)$ children.



| i | T_i | $\text{prev}(T_i)$ | $\text{prev}(T_{\text{pSA}[i]})$ | $T_{\text{pSA}[i]}$ | $\text{pSA}[i]$ | $L[i]$ | f_i | $\text{pBWT}[i]$ | $\text{pLF}(i)$ |
|-----|----------------|--------------------|----------------------------------|---------------------|-----------------|--------|-------|------------------|-----------------|
| 1 | AxyBzCxyzwAz\$ | A00B0C530A3\$ | 000A3\$A70B6C | xzwAz\$AxyBzC | 7 | C | | C | 11 |
| 2 | xyBzCxyzwAz\$A | 00B0C530A3\$A | 00A3\$A00B6C5 | zwAz\$AxyBzCx | 8 | x | 7 | 3 | 1 |
| 3 | yBzCxyzwAz\$Ax | 0B0C030A3\$A7 | 00B0C530A3\$A | xyBzCxyzwAz\$A | 2 | A | | A | 8 |
| 4 | BzCxyzwAz\$Axy | B0C030A3\$A70 | 0A0\$A00B6C53 | wAz\$AxyBzCxz | 9 | z | 3 | 2 | 2 |
| 5 | zCxyzwAz\$AxyB | 0C030A3\$A70B | 0B0C030A3\$A7 | yBzCxyzwAz\$Ax | 3 | x | 5 | 3 | 3 |
| 6 | CxyzwAz\$AxyBz | C000A3\$A70B6 | 0C030A3\$A70B | zCxyzwAz\$AxyB | 5 | B | | B | 10 |
| 7 | xzwAz\$AxyBzC | 000A3\$A70B6C | 0\$A00B5C530A | z\$AxyBzCxyzwA | 11 | A | | A | 9 |
| 8 | zwAz\$AxyBzCx | 00A3\$A00B6C5 | A00B0C530A3\$ | AxyBzCxyzwAz\$ | 1 | \$ | | \$ | 12 |
| 9 | wAz\$AxyBzCxz | 0A0\$A00B6C53 | A0\$A00B5C530 | Az\$AxyBzCxyzw | 10 | w | 12 | 4 | 4 |
| 10 | Az\$AxyBzCxyzw | A0\$A00B5C530 | B0C030A3\$A70 | BzCxyzwAz\$Axy | 4 | y | 12 | 4 | 5 |
| 11 | z\$AxyBzCxyzwA | 0\$A00B5C530A | C000A3\$A70B6 | CxyzwAz\$AxyBz | 6 | z | 3 | 2 | 6 |
| 12 | \$AxyBzCxyzwAz | \$A00B0C530A3 | \$A00B0C530A3 | \$AxyBzCxyzwAz | 12 | z | 6 | 3 | 7 |

Figure 2: The text is $T[1, 12] = AxyBzCxyzwAz\$$, where $\Sigma_s = \{A, B, C, \$\}$ and $\Sigma_p = \{w, x, y, z\}$

sorted matrix M . Clearly, the i th row is equal to $T_{\text{pSA}[i]}$. Moving forward, denote by f_i , the first occurrence of $L[i]$ in $T_{\text{pSA}[i]}$. The p-BWT of T , denoted by $\text{pBWT}[1, n]$, is defined as follows:

$$\text{pBWT}[i] = \begin{cases} L[i], & \text{if } L[i] \text{ is an s-character,} \\ \text{number of distinct p-characters in } T_{\text{pSA}[i]}[1, f_i], & \text{otherwise.} \end{cases}$$

In other words, when $L[i] \in \Sigma_s$, $\text{pBWT}[i] = T[\text{pSA}[i] - 1]$ (define $T[0] = T[n] = \$$) and when $L[i] \in \Sigma_p$, $\text{pBWT}[i]$ is the number of 0's in the f_i -long prefix of $\text{prev}(T_{\text{pSA}[i]})$. Thus, pBWT is a sequence of n characters over the set $\Sigma'' = \Sigma_s \cup \{1, 2, \dots, \sigma_p\}$ of size $\sigma_s + \sigma_p = \sigma$. See Figure 2 for an illustration of pSA and pBWT .

In order to represent pBWT in succinct space, we map each s-character in Σ'' to a unique integer in $[\sigma_p + 1, \sigma]$. Specifically, the i th smallest s-character will be denoted by $(i + \sigma_p)$. Moving forward, $\text{pBWT}[i] \in [1, \sigma_p]$ iff $L[i]$ is a p-character and $\text{pBWT}[i] \in [\sigma_p + 1, \sigma]$ iff $L[i]$ is a s-character. We summarize the relation between $\text{prev}(T_{\text{pSA}[i]})$ and $\text{prev}(T_{\text{pSA}[i-1]})$ below.

Observation 1 For any $1 \leq i \leq n$,

$$\text{prev}(T_{\text{pSA}[i-1]}) = \begin{cases} \text{pBWT}[i] \circ \text{prev}(T_{\text{pSA}[i]}[1, n-1]), & \text{if } \text{pBWT}[i] \in \Sigma_s, \\ 0 \circ \text{prev}(T_{\text{pSA}[i]}[1, f_i-1] \circ f_i \circ \text{prev}(T_{\text{pSA}[i]}[f_i+1, n-1]), & \text{otherwise.} \end{cases}$$

3.1 Parameterized Last-to-First Column Mapping

Based on the conceptual matrix M , the parameterized last-to-first column (pLF) mapping of i is the position at which the character at $L[i]$ lies in the first column of M . Specifically, $\text{pLF}(i) = \text{pSA}^{-1}[\text{pSA}[i] - 1]$. The significance of pLF mapping is summarized in Lemma 1.

Lemma 1 *Assume $\text{pLF}(\cdot)$ can be computed in t_{pLF} time. Then, for any parameter Δ , by using an additional $\mathcal{O}((n/\Delta) \log n)$ -bit structure, we can compute $\text{pSA}[\cdot]$ and $\text{pSA}^{-1}[\cdot]$ in $\mathcal{O}(\Delta \cdot t_{\text{pLF}})$ time.*

Proof: Define, $\text{pLF}^0(i) = i$ and $\text{pLF}^k(i) = \text{pLF}(\text{pLF}^{k-1}(i)) = \text{pSA}^{-1}[\text{pSA}[i] - k]$ for any integer $k > 0$. We maintain two Δ -sampled arrays, one each for pSA and pSA^{-1} . More specifically, we explicitly maintain $\text{pSA}[j]$ and $\text{pSA}^{-1}[j]$ if the value belongs to $\{1, 1 + \Delta, 1 + 2\Delta, 1 + 3\Delta, \dots, n\}$. The total space for each sampled array can be bounded by $\mathcal{O}((n/\Delta) \log n)$ bits. To find $\text{pSA}[i]$, repeatedly apply the $\text{pLF}(\cdot)$ operation (starting from i) until you obtain a j such that $\text{pSA}[j]$ has been explicitly stored. Suppose, the number of such operations is k . Then, $j = \text{pLF}^k(i) = \text{pSA}^{-1}[\text{pSA}[i] - k]$, which gives $\text{pSA}[i] = \text{pSA}[j] + k$. Since $k \leq \Delta$, $\text{pSA}[i]$ is computed in $\mathcal{O}(\Delta \cdot t_{\text{pLF}})$ time. To find $\text{pSA}^{-1}[i]$, find the smallest $j \geq i$ whose $\text{pSA}^{-1}[j]$ is explicitly stored. Then, $\text{pSA}^{-1}[i] = \text{pLF}^{j-i}(\text{pSA}^{-1}[j])$. As $j - i \leq \Delta$, the time is bounded by $\mathcal{O}(\Delta \cdot t_{\text{pLF}})$. ■

We have the following as a consequence of Lemma 1.

Corollary 1 *Assume $\text{pLF}(\cdot)$ can be computed in t_{pLF} time and $\text{pBWT}[\cdot]$ can be accessed in t_{pBWT} time. Then, for any parameter Δ , using an additional $\mathcal{O}((n/\Delta) \log n)$ -bit structure, $\text{prev}(\text{T}[x, y])$ can be extracted in $\mathcal{O}(\Delta \cdot t_{\text{pLF}} + d(t_{\text{pLF}} + t_{\text{pBWT}} + \log \sigma))$ time, where $d = y - x + 1$.*

Proof: Note that $\text{T}[n] = \$$ and by Convention 1, we have $\text{pSA}[n] = \text{pSA}^{-1}[n] = n$. We first retrieve the lexicographic position i of the circular suffix T_{y+1} by using $\text{pSA}^{-1}[y + 1]$, and then retrieve $\text{pBWT}[i]$. This is achieved in $\mathcal{O}(\Delta \cdot t_{\text{pLF}} + t_{\text{pBWT}})$ time. We now proceed to retrieve the character $\text{pBWT}[\text{pSA}^{-1}[y]]$. Note that although this can be directly achieved using pSA^{-1} , this will cost $\mathcal{O}(\Delta \cdot t_{\text{pLF}} + t_{\text{pBWT}})$ time. To achieve this in a more efficient way, we first find the position i' , such that the first character in suffix $\text{T}_{\text{pSA}[i']}$ is exactly the same character as $L[i]$. Therefore, $i' = \text{pLF}(i)$ and we obtain $\text{pBWT}[i']$. The time required for this step is $\mathcal{O}(t_{\text{pLF}} + t_{\text{pBWT}})$. We repeat this process until we obtain the lexicographic position i'' of the circular suffix T_{x+1} , and terminate by obtaining $\text{pBWT}[i'']$. At this point, we have obtained the following string

$$S = \text{pBWT}[\text{pLF}^d(i)] \circ \text{pBWT}[\text{pLF}^{d-1}(i)] \circ \dots \circ \text{pBWT}[\text{pLF}^2(i)] \circ \text{pBWT}[\text{pLF}(i)] \circ \text{pBWT}[i]$$

where $i = \text{pSA}^{-1}[y + 1]$. The time required to obtain the string is $\mathcal{O}(\Delta \cdot t_{\text{pLF}} + d(t_{\text{pLF}} + t_{\text{pBWT}}))$.

Recall that $\text{pBWT}[j]$ is $L[j]$ if $L[j]$ is an s-character; otherwise, it is number of distinct p-characters in $\text{T}_{\text{pSA}[j]}[1, f_j]$, where f_j is the first occurrence of $L[j]$ in $\text{T}_{\text{pSA}[j]}$. Therefore, the obtained string contains the desired s-characters of $\text{T}[x, y]$ at the correct position, and we are required to obtain the encoding corresponding to the p-characters in $\text{T}[x, y]$.

We scan the obtained string S from right to left. Suppose, we encounter a p-character in S at a position d' from left (which can be checked by comparing the $S[d']$ with σ_p). We find the number of distinct p-characters, say k' , in $S[d' + 1, d]$. This can be achieved in constant time⁴. (In case,

⁴ We maintain a bit-vector $B[1, \sigma_p]$, all values initialized to 0, and a counter C initialized to 0. Scan the text from right to left, and whenever a p-character c_p is encountered at position x , the required count at x is C . If $B[c_p] = 0$, then increment C by 1, set $B[c_p]$ to 1, and continue.

$d' = d$, then $k' = 0$.) Observe that if $S[d'] > k'$, then the p-character $T[x - 1 + d']$ does not appear in $T[x + d', y]$, and we replace $S[d']$ by a new p-character $c_{k'+1}$. Also, insert $c_{k'+1}$ into a *balanced binary search tree* keyed by d' . Otherwise, $T[x - 1 + d']$ is the $S[d']$ th smallest p-character in the tree. We retrieve this character from the tree and update the key to d' . Since, the tree has at most σ_p nodes, insertion, search, and key update operations are performed in $\mathcal{O}(\log \sigma_p)$ time. Therefore, this process takes $\mathcal{O}(d \log \sigma_p)$ time. Finally, we encode the string using $\text{prev}(\cdot)$ in $\mathcal{O}(d \log \sigma_p)$ time, which results in the claimed complexity. \blacksquare

To aid the reader's intuition for computing pLF mapping, we present Lemma 2, which shows how to compare the lexicographic rank of two encoded suffixes when prepended by their respective previous characters. This key concept is implemented in Sections 4 and 6 to arrive at Theorem 2.

Lemma 2 *Consider two suffixes i and j corresponding to the leaves ℓ_i and ℓ_j in pST. Then, $\text{pLF}(i)$ and $\text{pLF}(j)$ are related as follows:*

- (a) *If $L[i]$ is parameterized and $L[j]$ is static, then $\text{pLF}(i) < \text{pLF}(j)$.*
- (b) *If both $L[i]$ and $L[j]$ are static, then $\text{pLF}(i) < \text{pLF}(j)$ iff one of the following holds:*
 - $\text{pBWT}[i] < \text{pBWT}[j]$
 - $\text{pBWT}[i] = \text{pBWT}[j]$ and $i < j$
- (c) *Assume that both $L[i]$ and $L[j]$ are parameterized and $i < j$. Let u be the lowest common ancestor of ℓ_i and ℓ_j in pST, and z be the number of 0's in the string $\text{path}(u)$. Then,*
 - (1) *If $\text{pBWT}[i], \text{pBWT}[j] \leq z$, then $\text{pLF}(i) < \text{pLF}(j)$ iff $\text{pBWT}[i] \geq \text{pBWT}[j]$.*
 - (2) *If $\text{pBWT}[i] \leq z < \text{pBWT}[j]$, then $\text{pLF}(i) > \text{pLF}(j)$.*
 - (3) *If $\text{pBWT}[i] > z \geq \text{pBWT}[j]$, then $\text{pLF}(i) < \text{pLF}(j)$.*
 - (4) *If $\text{pBWT}[i], \text{pBWT}[j] > z$, then $\text{pLF}(i) > \text{pLF}(j)$ iff all of the following are true:*
 - $\text{pBWT}[i] = z + 1$.
 - *the leading character on the u to ℓ_i path is 0.*
 - *the leading character on the u to ℓ_j path is not an s-character.*

Proof: (a) and (b): Follows immediately from Convention 1 and Observation 1.

(c) Recall that f_i and f_j are the first occurrences of the characters $L[i]$ and $L[j]$ in the circular suffixes $T_{\text{pSA}[i]}$ and $T_{\text{pSA}[j]}$ respectively. Let $d = |\text{path}(u)|$. Clearly, the conditions (1)–(4) can be written as: (1) Both $f_i, f_j \leq d$, (2) $f_i \leq d$ and $f_j > d$, (3) $f_i > d$ and $f_j \leq d$, and (4) Both $f_i, f_j > d$.

Then the claims (1)–(3) are immediate from Observation 1 and Convention 1. For proving (4), first observe that if $T_{\text{pSA}[j]}[d + 1]$ is an s-character, then $T_{\text{pSA}[j]-1}[d + 2] > T_{\text{pSA}[i]-1}[d + 2]$, and $\text{pLF}(i) < \text{pLF}(j)$. So, assume otherwise. Let e_i and e_j be the $(d + 1)$ th characters of $\text{prev}(T_{\text{pSA}[i]})$ and $\text{prev}(T_{\text{pSA}[j]})$ respectively. Since, the suffixes i and j separates after u , $f_i \neq f_j$. Also, $i < j$ implies $0 \leq e_i < e_j \leq d$. Note that if $\text{pBWT}[i] = z + 1$ and $e_i = 0$, then $L[i] = T_{\text{pSA}[i]}[d + 1]$ i.e., $f_i = d + 1$, and $\text{prev}(T_{\text{pSA}[i]-1})[d + 2] = d + 1 > e_j = \text{prev}(T_{\text{pSA}[j]-1})[d + 2]$. Otherwise, $\text{prev}(T_{\text{pSA}[i]-1})[d + 2] = e_i < e_j \leq \text{prev}(T_{\text{pSA}[j]-1})[d + 2]$. \blacksquare

Theorem 2 *We can compute $\text{pLF}(i)$ as follows:*

- (a) in $\mathcal{O}(\log \sigma)$ time using $\mathcal{O}(n \log \sigma)$ bits.
- (b) in $\mathcal{O}(\log \sigma \log \log \sigma)$ time using $n \log \sigma + \mathcal{O}(n)$ bits.

4 Implementing pLF Mapping in Compact Space

We prove Theorem 2a in this section.

4.1 Data Structure Toolkit

The following form the key components of our data structure.

4.1.1 Wavelet Tree over pBWT

Grossi, Gupta, and Vitter [GGV03] introduced the wavelet tree (WT) data structure, which generalizes the well-known rank and select queries over bit-vectors⁵. Specifically, given an array A over an alphabet Σ , by maintaining a data structure of size $|A| \log |\Sigma| + o(|A| \log |\Sigma|)$ bits, the following queries can be supported in $\mathcal{O}(\log |\Sigma|)$ time [GGV03, GNP12, Nav14]:

- $A[i]$.
- $\text{rank}_A(i, x)$ = the number of occurrences of x in $A[1, i]$.
- $\text{select}_A(i, x)$ = the i th occurrence of x in $A[1, n]$.
- $\text{predecessor}_A(i, x) = j$, the rightmost position before i such that $A[j] < x$.
- $\text{rangeCount}_A(i, j, x, y)$ = number of elements in $A[i, j]$ that are at least x and at most y .
- $\text{rangeNextVal}_A(i, j, x, k)$ = the k th smallest value in $A[i, j]$ that is greater than x .
- $\text{rangePrevVal}_A(i, j, x, k)$ = the k th largest value in $A[i, j]$ that is smaller than x .

We drop the subscript A when the context is clear. The pBWT is a string of length n over an alphabet set $\Sigma'' = \Sigma_s \cup \{1, 2, \dots, \sigma_p\}$ of size $\sigma = \sigma_s + \sigma_p$. By maintaining a WT over pBWT in $n \log \sigma + o(n \log \sigma)$ bits, we can support the above operations over the pBWT. As noted by Navarro [Nav14], we can apply the technique of Golynski et al. [GGG⁺07] to reduce the redundancy of $o(n \log \sigma)$ bits to $o(n)$ bits. The time to answer any of the above queries remains unaffected.

4.1.2 Succinct representation of the topology of pST

We rely on the following result of Navarro and Sadakane [NS14, SN10]. Any tree having m nodes can be represented in $2m + o(m)$ bits, such that if each node is labeled by its pre-order rank in the tree, the following operations can be supported in constant time (note that $m < 2n$ in our case):

- $\text{pre-order}(u)$ = the pre-order rank of node u .
- $\text{parent}(u)$ = the parent of node u .

⁵ Given a bit-vector B and $c \in \{0, 1\}$, $\text{rank}(i, c) = |\{j \mid j \leq i \text{ and } B[j] = c\}|$ and $\text{select}(i, c) = \min\{j \mid \text{rank}(j, c) = i\}$

- $\text{nodeDepth}(u)$ = the number of edges on the path from root to u .
- $\text{child}(u, q)$ = the q th leftmost child of node u .
- $\text{lca}(u, v)$ = the lowest common ancestor (LCA) of two nodes u and v .
- $\text{lmostLeaf}(u)/\text{rmostLeaf}(u)$ = the leftmost/rightmost leaf of the subtree rooted at u .
- $\text{levelAncestor}(u, D)$ = the ancestor of u such that $\text{nodeDepth}(u) = D$.

The representation also allows us to find the pre-order rank of the i th leftmost leaf in $\mathcal{O}(1)$ time. Thus, for a particular suffix i , we can locate its corresponding leaf ℓ_i in this representation (i.e., find the pre-order rank of ℓ_i) in $\mathcal{O}(1)$ time. Given a leaf ℓ , in $\mathcal{O}(1)$ time, we can also find the number of suffixes which appear before the suffix corresponding to ℓ . Moving forward, we will use ℓ_i to denote the leaf corresponding to the i th lexicographically smallest prev-encoded suffix.

4.1.3 Marked and Prime Nodes

We use the following scheme of Hon et al. [HSTV14]. Identify certain nodes in a tree (pST in our case) as *marked nodes*. Let g be a parameter called grouping factor. Starting from the leftmost leaf in pST, we combine every g leaves together to form a group. In particular, the leaves ℓ_1 through ℓ_g form the first group, ℓ_{g+1} through ℓ_{2g} form the second, and so on. We mark the LCA of the first and last leaves of every group. Moreover, for any two consecutive marked nodes in pre-order, we mark their LCA as well, and continue recursively. Note that the root node is marked. For every marked node u^* , we also identify a *prime node*, which is the child of the lowest marked ancestor u' of u^* that lies on the path from u' to u^* . If $\text{parent}(u^*)$ is marked then u^* 's prime node is u^* itself.

Fact 2 ([HSTV14]) *The following properties are ensured by the above marking scheme.*

- (1) *The number of marked nodes and prime nodes is $\mathcal{O}(n/g)$.*
- (2) *For any (prime) node u' , the highest marked node u^* in its subtree is unique. Let $\text{leaf}(v)$ be the set of leaves in the subtree of a node v . Then, $\text{leaf}(u') \setminus \text{leaf}(u^*)$ is of size at most $2g$. Also, for a lowest marked node u , $|\text{leaf}(u)| \leq 2g$.*
- (3) *The path from a (marked) node to its nearest marked ancestor contains at most g nodes.*
- (4) *Given a node, in $\mathcal{O}(1)$ time, we can verify whether it has a marked descendant or not. We can also find the highest marked descendant (if any) in $\mathcal{O}(1)$ time⁶.*
- (5) *Given a node, in $\mathcal{O}(1)$ time, we can find its lowest marked ancestor and its lowest prime ancestor (if any) using an $\mathcal{O}(n)$ -bit structure⁷.*

⁶ Let i and j be two multiples of g such that $j > i \geq 0$. Clearly, for any such i, j pair, if the subtree of u contains leaves ℓ_{i+1} and ℓ_j , then there u has a marked descendant, otherwise not. The LCA of such a maximal pair of leaves ℓ_{i+1} and ℓ_j gives the highest marked descendant.

⁷ Maintain a bit-vector B such that $B[i] = 1$ iff the node with pre-order rank i is a marked node. Additionally, maintain a rank-select data structure over this bit-vector. Now given a query node with pre-order rank j , using these data structures, in constant time, find the largest $j' < j$ such that $B[j'] = 1$ and the smallest $j'' > j$ such that $B[j''] = 1$. Note that one of j' or j'' must exist. If j' is an ancestor of j (which can be verified by comparing the ranges of the leaves in their subtree), then j' is the desired lowest marked ancestor. In case, j' or j'' does

4.2 ZeroDepth and ZeroNode

For a node u , $\text{zeroDepth}(u)$ is the number of 0's in $\text{path}(u)$. For a leaf ℓ_i with $\text{pBWT}[i] \in [1, \sigma_p]$, $\text{zeroNode}(\ell_i)$ is the highest node z on the root to ℓ_i path such that $\text{zeroDepth}(z) \geq \text{pBWT}[i]$. Thus, z is the locus of $\text{path}(\ell_i)[1, f_i]$. Note that z necessarily exists as $\text{zeroDepth}(\ell_i) \geq \text{pBWT}[i]$. Moving forward, whenever we refer to $\text{zeroNode}(\ell_i)$, we assume $\text{pBWT}[i] \in [1, \sigma_p]$.

Lemma 3 *We can compute $\text{zeroNode}(\ell_i)$ in $\mathcal{O}(\log \sigma)$ time using an $\mathcal{O}(n \log \sigma)$ -bit structure.*

Proof: We find the lowest node u on the root to ℓ_i path that satisfies $\text{zeroDepth}(u) < \text{pBWT}[i]$. Clearly, $\text{zeroNode}(\ell_i)$ is the child of u on this path, and is given by $\text{levelAncestor}(\ell_i, \text{nodeDepth}(u) + 1)$. Since the zeroDepth of the root node is 0, u necessarily exists.

Maintain an array D such that $D[k]$ equals zeroDepth of the node with pre-order rank k . Instead of maintaining D explicitly, we maintain a wavelet tree over it. The space required is $2n \log \sigma + o(n)$ bits. We locate the node, say v , with pre-order rank $\text{predecessor}_D(j, \text{pBWT}[i])$, where j is the pre-order rank of ℓ_i . Then, u is given by $\text{lca}(\ell_i, v)$. The time required is $\mathcal{O}(\log \sigma)$.

To see the correctness first observe that for any node x , $\text{zeroDepth}(x) \geq \text{zeroDepth}(\text{parent}(x))$. Thus, $\text{zeroDepth}(\text{lca}(\ell_i, v)) \leq \text{zeroDepth}(v) < \text{pBWT}[i]$. If $\text{lca}(\ell_i, v)$ is not the desired node, then its child u' on the path to ℓ_i satisfies $\text{zeroDepth}(u') < \text{pBWT}[i]$. But then u' appears after v and before ℓ_i in pre-order, a contradiction. ■

We remark that the following additional functionalities: $\text{leafLeadChar}(\cdot)$, $\text{fSum}(\cdot)$ and $\text{pCount}(\cdot)$ will be defined later. Each of these can be computed in $\mathcal{O}(1)$ time using an $\mathcal{O}(n)$ -bit structure. In summary, the total space required by the above data structures is $\mathcal{O}(n \log \sigma)$ bits.

4.3 Computing $\text{pLF}(i)$ when $\text{pBWT}[i] \in [\sigma_p + 1, \sigma]$

In this case, $L[i] = \text{pBWT}[i]$ is an s-character. Using Lemma 2, we conclude that $\text{pLF}(i) > \text{pLF}(j)$ iff either $j \in [1, n]$ and $\text{pBWT}[j] < \text{pBWT}[i]$, or $j \in [1, i - 1]$ and $\text{pBWT}[i] = \text{pBWT}[j]$. Then, $\text{pLF}(i) = 1 + \text{rangeCount}(1, n, 1, c - 1) + \text{rangeCount}(1, i - 1, c, c)$, where $c = \text{pBWT}[i]$.

4.4 Computing $\text{pLF}(i)$ when $\text{pBWT}[i] \in [1, \sigma_p]$

In this case, $L[i]$ is a p-character. Let $z = \text{zeroNode}(\ell_i)$ and $v = \text{parent}(z)$. Then, $f_i = (|\text{path}(v)| + 1)$ if the leading character on the edge from v to z is 0 and $\text{pBWT}[i] = (\text{zeroDepth}(v) + 1)$; otherwise, $f_i > (|\text{path}(v)| + 1)$. For a leaf ℓ_j in pST , $\text{leafLeadChar}(j)$ is a boolean variable, which is 0 iff $f_j = (|\text{path}(\text{parent}(\text{zeroNode}(\ell_j)))| + 1)$. Using this information, in constant time, we can determine which of the following two cases the suffix corresponding to ℓ_i satisfies.

4.4.1 Case 1 ($f_i = |\text{path}(v)| + 1$)

See Figure 3a for an illustration. Let w be the parent of v . Note that in this case z is the leftmost child of v . We partition the leaves into four sets \mathcal{S}_1 , \mathcal{S}_2 , \mathcal{S}_3 , and \mathcal{S}_4 :

not exist, we take them to be the root node. Consider the nodes $\text{lca}(j, j')$, $\text{lca}(j, j'')$, and $\text{lca}(j', j'')$. The desired lowest marked ancestor u is the one among these three which is marked and has maximum node depth. Once u is retrieved, we check whether it has marked descendant, or not. If it has one, the lowest prime ancestor is given by $\text{levelAncestor}(j, \text{nodeDepth}(u) + 1)$.




Figure 3: Illustration of various suffix ranges when the suffix $T_{\text{pSA}[i]}$ is preceded by a p-character

- \mathcal{S}_1 : leaves to the left of the subtree of v .
- \mathcal{S}_2 : leaves in the subtree of z .
- \mathcal{S}_3 : leaves to the right of the subtree of v .
- \mathcal{S}_4 : leaves in the subtree of v but not of z .

In case, v is the root node r , we take w to be r as well. Consequently, $\mathcal{S}_1 = \mathcal{S}_3 = \emptyset$.

4.4.2 Case 2 ($f_i > |\text{path}(v)| + 1$)

See Figure 3b for an illustration. We partition the leaves into three sets \mathcal{S}_1 , \mathcal{S}_2 , and \mathcal{S}_3 :

- \mathcal{S}_1 (resp. \mathcal{S}_3): leaves to the left (resp. right) of the subtree of z .
- \mathcal{S}_2 : leaves in the subtree of z .

4.4.3 Computing $\text{LF}(i)$

We first compute $z = \text{zeroNode}(\ell_i)$ using Lemma 3, and then locate $v = \text{parent}(z)$. Using $\text{leafLeadChar}(i)$ and the $\text{lmostLeaf}(\cdot)/\text{rmostLeaf}(\cdot)$ tree operations, we find the desired ranges.

We shall use $[L_x, R_x]$ to denote the range of leaves in the subtree of any node x . In order to compute $\text{pLF}(i)$, we first compute N_1 , N_2 , and N_3 , which are respectively the number of leaves ℓ_j in ranges \mathcal{S}_1 , \mathcal{S}_2 , and \mathcal{S}_3 such that $\text{pLF}(j) \leq \text{pLF}(i)$. Likewise, we compute N_4 (w.r.t \mathcal{S}_4) if we are in the first case. Then, $\text{pLF}(i) = N_1 + N_2 + N_3 + N_4$. We present the detailed description below.

- **Computing N_1 :** For any leaf $\ell_j \in \mathcal{S}_1$, $\text{pLF}(j) < \text{pLF}(i)$ iff $f_j > 1 + |\text{path}(\text{lca}(z, \ell_j))|$ and $L[j] \in \Sigma_p$. Therefore, N_1 is the number of leaves ℓ_j , $L[j] \in \Sigma_p$, which comes before z in pre-order with $f_j > 1 + |\text{path}(\text{lca}(z, \ell_j))|$.

Define, $\text{fCount}(x)$ of a node x as the number of leaves ℓ_j in x 's subtree such that $|\text{path}(y)| + 2 \leq f_j \leq |\text{path}(x)| + 1$, where $y = \text{parent}(x)$. If x is the root node, then $\text{fCount}(x) = 0$.

Define $\text{fSum}(x)$ of a node x as the sum of $\text{fCount}(y)$ of all nodes y which come before x in pre-order and are not ancestors of x .

By this definition, $N_1 = \text{fSum}(z)$ is computed as follows (proof deferred to Section 4.4.4).

Lemma 4 *By maintaining an $\mathcal{O}(n)$ -bit structure, we can compute $\text{fSum}(x)$ in $\mathcal{O}(1)$ time.*

- **Computing N_2 :** Note that for any leaf $\ell_j \in \mathcal{S}_2$, $\text{pLF}(j) \leq \text{pLF}(i)$ iff $L[j] \in \Sigma_p$ and either $f_j > f_i$ or $f_j = f_i$ and $j \leq i$. Therefore, N_2 is the number of leaves ℓ_j in \mathcal{S}_2 which satisfy one of the following conditions: (a) $\text{pBWT}[i] < \text{pBWT}[j] \leq \sigma_p$, or (b) $\text{pBWT}[i] = \text{pBWT}[j]$ and $j \leq i$. Then, $N_2 = \text{rangeCount}(L_z, R_z, c + 1, \sigma_p) + \text{rangeCount}(L_z, i, c, c)$, where $c = \text{pBWT}[i]$.
- **Computing N_3 :** For any leaf $\ell_j \in \mathcal{S}_3$, $\text{pLF}(j) > \text{pLF}(i)$. Therefore, $N_3 = 0$.
- **Computing N_4 :** Note that $\text{pBWT}[i]$ is same as $(\text{zeroDepth}(v) + 1)$. Consider a leaf $\ell_j \in \mathcal{S}_4$ with $L[j] \in \Sigma_p$. Since the suffix j deviates from the suffix i at the node v , we have $f_j \neq f_i$. Therefore, $\text{pLF}(j) < \text{pLF}(i)$ iff $f_j > f_i$, and the leading character on the path from v to ℓ_j is not an s-character. For a node x , $\text{pCount}(x)$ is the number of children y of x such that the leading character from x to y is not an s-character. Note that $\sum_x \text{pCount}(x) = \mathcal{O}(n)$. Therefore, we encode $\text{pCount}(\cdot)$ of all nodes in $\mathcal{O}(n)$ bits using unary encoding, such that $\text{pCount}(x)$ can be retrieved in constant time⁸. Let u be the $\text{pCount}(v)$ th child of v . Then, N_4 is the number of leaves ℓ_j in \mathcal{S}_4 such that $j \leq R_u$ and $\sigma_p \geq \text{pBWT}[j] \geq \text{pBWT}[i]$ i.e., $N_4 = \text{rangeCount}(R_z + 1, R_u, \text{pBWT}[i], \sigma_p)$.

We summarize the LF mapping procedure in Algorithm 2. Once $\text{zeroDepth}(\ell_i)$ is known, N_1 is computed in $\mathcal{O}(1)$ time, and both N_2 and N_4 are computed in $\mathcal{O}(\log \sigma)$ time. Combining these results with Lemma 3, we arrive at Theorem 2a.

4.4.4 Proof of Lemma 4

We first present the underlying intuition. Since $\sum_x \text{fCount}(x) \leq n$, $\text{fCount}(\cdot)$ of all nodes can be encoded in $\mathcal{O}(n)$ bits using unary encoding. We can decode $\text{fCount}(\cdot)$ of any particular node in constant time. We carefully select a set of $\mathcal{O}(\frac{n}{\log n})$ nodes, and explicitly storing their $\text{fSum}(\cdot)$ values in $\mathcal{O}(n)$ bits. The $\text{fCount}(\cdot)$ of the remaining nodes can be encoded efficiently using the Method of the Four Russians [Gus97]. Now, we present the details.

We mark nodes in the pST w.r.t grouping factor $g = \lceil \log n \rceil / \lambda$, where λ is a constant to be determined later. For every prime node v' and a lowest marked node v^* , we store $\text{fSum}(v')$ and $\text{fSum}(v^*)$. For every marked node w^* , we store $\Phi(w^*)$ i.e., the sum of $\text{fCount}(w)$ for every w in the subtree rooted at w^* . According to Fact 2, the number of prime nodes and marked nodes is $\mathcal{O}(n/g)$. Therefore, the space needed is $\mathcal{O}(\frac{n}{g} \log n) = \mathcal{O}(n)$ bits⁹.

⁸ Create a binary string S as follows. For each node u in a pre-order traversal of pST, append to S a 0 followed by $\text{pCount}(u)$ 1s. Append a 0 at the end, and maintain a rank-select structure over S . Then $\text{pCount}(u)$ for a node u , having pre-order rank k , is the number of 1s between the k th 0 and the $(k+1)$ th 0. The value is found by using two select operations followed by two rank operations.

⁹ Maintain a bit-vector B such that $B[k] = 1$ iff the node with pre-order rank k is a prime node w.r.t g . Keep a rank-select structure over B . In an adjoining array F , store the fSum value of each prime node i.e., $F[k]$ stores

Algorithm 2 computes $\text{pLF}(i)$

```

1:  $c \leftarrow \text{pBWT}[i]$ 
2: if  $(c > \sigma_p)$  then
3:    $\text{pLF}(i) \leftarrow 1 + \text{rangeCount}(1, n, 1, c - 1) + \text{rangeCount}(1, i - 1, c, c)$ 
4: else
5:    $z \leftarrow \text{zeroNode}(\ell_i), v \leftarrow \text{parent}(z)$ 
6:    $N_1 \leftarrow \text{fSum}(z)$ 
7:    $L_z \leftarrow \text{lmostLeaf}(z), R_z \leftarrow \text{rmostLeaf}(z)$ 
8:    $N_2 \leftarrow \text{rangeCount}(L_z, R_z, c + 1, \sigma_p) + \text{rangeCount}(L_z, i, c, c)$ 
9:    $N_3 \leftarrow 0$ 
10:  if  $(\text{leafLeadChar}(i) \text{ is } 0)$  then
11:     $u \leftarrow \text{child}(v, \text{pCount}(v))$ 
12:     $N_4 \leftarrow \text{rangeCount}(R_z + 1, \text{rmostLeaf}(u), c, \sigma_p)$ 
13:  end if
14:   $\text{pLF}(i) \leftarrow N_1 + N_2 + N_4$ 
15: end if

```

Consider a prime node u' and let u^* be its (unique) highest marked descendant. Let the sequence of nodes on the path from u' to u^* be $u' = u_0, u_1, u_2, \dots, u_k, u_{k+1} = u^*$. We store the following string corresponding to u' :

$$B_{u'} = 1^{n_1} \circ 0 \circ 1^{n_2} \circ 0 \circ \dots \circ 1^{n_k} \circ 0$$

Here, n_i is the number of leaves j in the subtree of u^* such that

$$|\text{path}(u_{i-1})| + 2 \leq f_j \leq |\text{path}(u_i)| + 1$$

Maintain a **rank-select** structure for $B_{u'}$, and associate it with the prime node u' . For a node u_i , $1 \leq i < k$, we define $\Gamma(u_i) = \sum_{i < i' \leq k} n_{i'}$. Define $\Gamma(u_k) = 0$. Observe that $\Gamma(u_i) = \text{rank}_{B_{u'}}(|B_{u'}|, 1) - \text{rank}_{B_{u'}}(i', 1)$, where $i' = \text{select}_{B_{u'}}(i, 0)$, is computed in $\mathcal{O}(1)$ time. Note that the sum of $|B_{u'}|$ over all prime nodes u' is at most $2n$. Therefore, the space needed is $\mathcal{O}(n)$ bits.

For any prime node u' and its highest marked descendant u^* , the number of nodes in the subtree \mathcal{T} of u' but not of u^* is at most $4g$ (since **pST** is a compacted trie and according to Fact 2). Therefore, we can encode the subtree \mathcal{T} in at most $8g$ bits [SN10]. Likewise, the subtree rooted at a lowest marked node can be encoded in at most $8g$ bits. Now, consider the *local pre-order rank* of any node v in \mathcal{T} . Given the pre-order rank of a node in **pST**, we can easily find its local pre-order rank in constant time. For every node $v \in \mathcal{T}$, we classify them into the following two categories (see Figure 4 for an illustration):

1. v either lies on or to the left of the $u'-u^*$ path i.e., v comes before u^* in pre-order
2. v lies to the right of the $u'-u^*$ path i.e., v comes after u^* in pre-order

the **fSum** value of the prime node corresponding to the k th 1-bit in B . Then, given the pre-order rank of a prime node, we can find its **fSum** value in $\mathcal{O}(1)$ time using **rank-select** operations on B to locate the desired position in F . Likewise, **fSum** of a lowest marked node and Φ of a marked node can be retrieved in $\mathcal{O}(1)$ time. The space needed for each structure is $\mathcal{O}(n)$ bits.

the space for all such encoding is again $o(n)$ bits.

Summarizing, the total space needed by the data structure is $\mathcal{O}(n)$ bits. Now, we concentrate on how to find $\text{fSum}(v_q)$ for a query node v_q .

If v_q is a prime node or a lowest marked node, then we can retrieve $\text{fSum}(v_q)$ in $\mathcal{O}(1)$ time. So, assume otherwise. We first determine whether it is in the subtree of a lowest marked ancestor, or not. (This is checked by determining v_q 's lowest marked ancestor and then checking whether that has a marked descendant, or not.) We have the following two cases.

- If v_q is not in the subtree of a lowest marked node, then we locate the lowest prime ancestor v'_q of v_q . Now, we locate the unique highest marked node v_q^* corresponding to v'_q . By comparing the pre-order ranks of v_q and v_q^* , we determine whether v_q is Category 1 or 2.

For Category 1, $\text{fSum}(v_q)$ is obtained using $\Psi_1(v_q)$ (stored in the table at v'_q) and $\text{fSum}(v'_q)$ i.e., $\text{fSum}(v_q) = \text{fSum}(v'_q) + \Psi_1(v_q)$.

For Category 2, we have $\text{fSum}(v_q) = \Gamma(\text{lca}(v_q, v_q^*)) + \text{fSum}(v'_q) + \Phi(v_q^*) + \Psi_2(v_q)$.

- Otherwise, we locate the lowest marked ancestor w_q^* of v_q . Using $\Psi(v_q)$ and $\text{fSum}(w_q^*)$, we compute $\text{fSum}(v_q)$.

We carry out a constant number of operations in either case, each taking $\mathcal{O}(1)$ time. This completes the proof of Lemma 4.

4.5 A Note on Construction

Given the p-suffix tree, every component of our index can be constructed in $\mathcal{O}(n \log \sigma)$ time using $\mathcal{O}(n \log n)$ bits of working space. Therefore, by first creating pST using Kosaraju's algorithm [Kos95], we have an $\mathcal{O}(n \log \sigma)$ time and $\mathcal{O}(n \log n)$ bit construction algorithm.

5 Pattern Matching via Backward Search

We use an adaptation of the backward search algorithm in the FM-index [FM00]. In particular, given a proper suffix Q of P , assume that we know the suffix range $[\text{sp}_1, \text{ep}_1]$ of $\text{prev}(Q)$. Our task is to find the suffix range $[\text{sp}_2, \text{ep}_2]$ of $\text{prev}(c \circ Q)$, where c is the character previous to Q in P .

If c is static i.e., $c \in [\sigma_p + 1, \sigma]$, then $\text{prev}(c \circ Q) = c \circ \text{prev}(Q)$. The backward search in this case is similar to that in FM-index. Specifically, $\text{sp}_2 = 1 + \text{rangeCount}(1, n, 1, c - 1) + \text{rangeCount}(1, \text{sp}_1 - 1, c, c)$ and $\text{ep}_2 = \text{rangeCount}(1, n, 1, c - 1) + \text{rangeCount}(1, \text{ep}_1, c, c)$. The time required is $\mathcal{O}(\log \sigma)$.

Now, we consider the scenario when c is parameterized. The idea is to compute the size of the suffix range of $\text{prev}(c \circ Q)$ i.e., $(\text{ep}_2 - \text{sp}_2 + 1)$, and then locate the suffix sp_2 . By maintaining a bit-vector $B[1, \sigma_p]$, in $\mathcal{O}(|P|)$ time, we first identify all positions j , where $P[j] \in \Sigma_p$ is not in $P[j+1, |P|]$. We handle the following two cases separately, and summarize the procedure in Algorithm 3.

5.1 Case 1 (c does not appear in Q)

Note that $\text{pLF}(i) \in [\text{sp}_2, \text{ep}_2]$ iff $i \in [\text{sp}_1, \text{ep}_1]$, $L[i]$ is a p-character and $f_i > |Q|$. This holds iff $i \in [\text{sp}_1, \text{ep}_1]$ and $\text{pBWT}[i] \in [d + 1, \sigma_p]$. Here, d is the number of distinct p-characters in Q , which

can be obtained in $\mathcal{O}(1)$ time by initially pre-processing P in $\mathcal{O}(|P|)$ time¹¹.

Then, $(\text{ep}_2 - \text{sp}_2 + 1) = \text{rangeCount}(\text{sp}_1, \text{ep}_1, d + 1, \sigma_p)$ is computed in time $\mathcal{O}(\log \sigma)$. Now, $\text{pLF}(i) < \text{sp}_2$ iff $i < \text{sp}_1$, $L[i] \in \Sigma_p$, and $f_i > 1 + |\text{path}(\text{lca}(u, \ell_i))|$, where $u = \text{lca}(\ell_{\text{sp}_1}, \ell_{\text{ep}_1})$. Therefore, we can compute $\text{sp}_2 = 1 + \text{fSum}(u)$ in constant time (refer to Lemma 4).

5.2 Case 2 (c appears in Q)

Note that $\text{pLF}(i) \in [\text{sp}_2, \text{ep}_2]$ iff $i \in [\text{sp}_1, \text{ep}_1]$, $L[i]$ is a p-character, and f_i is the same as the first occurrence of c in Q . This holds iff $i \in [\text{sp}_1, \text{ep}_1]$ and $\text{pBWT}[i] = d$. Here, d is the number of distinct p-characters in Q until (and including) the first occurrence of c . We can compute d in constant time by initially pre-processing P in $\mathcal{O}(|P| \log \sigma)$ time¹².

Then, $(\text{ep}_2 - \text{sp}_2 + 1) = \text{rangeCount}(\text{sp}_1, \text{ep}_1, d, d)$ is computed in time $\mathcal{O}(\log \sigma)$. Consider $i, j \in [\text{sp}_1, \text{ep}_1]$ such that $i < j$ and $\text{pLF}(i), \text{pLF}(j) \in [\text{sp}_2, \text{ep}_2]$. Now, both f_i and f_j equals the first occurrence of c in Q . Based on Observation 1, we conclude that $\text{pLF}(i) < \text{pLF}(j)$. Let $i_{\min} = \min\{j \mid j \in [\text{sp}_1, \text{ep}_1] \text{ and } \text{pBWT}[j] = d\}$. Note i_{\min} can be computed using the WT over pBWT in $\mathcal{O}(\log \sigma)$ time. Then, $\text{sp}_2 = \text{pLF}(i_{\min})$ is computed in $\mathcal{O}(\log \sigma)$ time (refer to Theorem 2a).

Summarizing, we find the suffix range of $\text{prev}(P)$ in $\mathcal{O}(|P| \log \sigma)$ time. Choose $\Delta = \lceil \log_\sigma n \rceil$ in Lemma 1. According to Theorem 2a, we have $t_{\text{pLF}} = \mathcal{O}(\log \sigma)$. Theorem 1a is now immediate.

6 Squeezing to Succinct Space

In this section, we prove Theorem 2b, and then use it to prove Theorem 1b. To achieve the space bound of Theorem 2b, we first need an alternative way (to that of Lemma 3) for computing zeroNode . To this end, we present the following lemma.

Lemma 5 *We can find $\text{zeroNode}(\ell_i)$ in $\mathcal{O}(\log \sigma \log \log \sigma)$ time using the WT over pBWT and an additional $\mathcal{O}(n)$ -bit structure.*

Based on Section 4 and the above lemma, Theorem 2b is immediate. Therefore, we can find the suffix range of $\text{prev}(P)$ in $\mathcal{O}(|P| \log \sigma \log \log \sigma)$ time (see Section 5). The additional $\mathcal{O}(n \log \sigma)$ space (above the $n \log \sigma$ space needed for the WT over pBWT) in Theorem 1a is due to the choice of $\Delta = \lceil \log_\sigma n \rceil$ in Lemma 1. We can improve this space to $\mathcal{O}(n)$ bits by choosing $\Delta = \lceil \log n \rceil$. Reporting a suffix array value requires $\mathcal{O}(\log n \log \sigma \log \log \sigma)$ time. Theorem 1b follows immediately.

The rest of this section is dedicated for proving Lemma 5. Before, getting into the details, we first present the key intuition. Let u_1 and u_2 be two consecutive nodes on the path from root to ℓ_i , and α be the number of leaves ℓ_j in the subtree of u_2 such that f_j lies on the edge (u_1, u_2) . Then, if we know $\text{zeroDepth}(u_1)$ or the largest $\text{pBWT}[j] \leq \text{zeroDepth}(u_1)$ for a leaf ℓ_j in the subtree of u_1 , we can find β i.e., the largest $\text{pBWT}[j'] \leq \text{zeroDepth}(u_2)$ value for a leaf $\ell_{j'}$ under the subtree of u_2

¹¹ Maintain a bit-vector $B[1, \sigma_p]$, all values initialized to 0, and a counter C initialized to 0. Scan the text from right to left. If a p-character c_p is encountered at position x , the required count at x is C . If $B[c_p] = 0$, then increment C by 1, set $B[c_p]$ to 1, and continue.

¹² We maintain an array $F[1, \sigma_p]$, all values initialized to 0, and a balanced binary search tree (initially empty). Scan the string from right to left, and when a p-character c_p is encountered at a position x , check $F[c_p]$. If $F[c_p] = 0$, then insert c_p in the tree keyed by x . Otherwise, the count at position x is the number of node in the tree with key at most $F[c_p]$. Update $F[c_p]$ and the key of c_p in the tree to x . Since the size of the tree is $\mathcal{O}(\sigma_p)$, search, insertion, and key update time is $\mathcal{O}(\log \sigma)$.

Algorithm 3 computes Suffix Range of $\text{prev}(P[1, p])$

```

1:  $c \leftarrow P[p]$ ,  $i \leftarrow p$ ,  $P \leftarrow \text{prev}(P)$ 
2: if ( $c > \sigma_p$ ) then
3:    $\text{sp} \leftarrow 1 + \text{rangeCount}(1, n, 1, c - 1)$ 
4:    $\text{ep} \leftarrow \text{rangeCount}(1, n, 1, c)$ 
5: else
6:    $\text{sp} \leftarrow 1$ 
7:    $\text{ep} \leftarrow 1 + \text{rangeCount}(1, n, 1, \sigma_p)$ 
8: end if
9: while ( $\text{sp} \leq \text{ep}$  and  $i \geq 2$ ) do
10:   $c \leftarrow P[i - 1]$ 
11:  if ( $c > \sigma_p$ ) then
12:     $\text{sp} \leftarrow 1 + \text{rangeCount}(1, n, 1, c - 1) + \text{rangeCount}(1, \text{sp} - 1, c, c)$ 
13:     $\text{ep} \leftarrow \text{rangeCount}(1, n, 1, c - 1) + \text{rangeCount}(1, \text{ep}, c, c)$ 
14:  else if ( $c \notin P[i, p]$ ) then
15:     $d \leftarrow$  number of distinct p-characters in  $P[i, p]$ 
16:     $\text{sp}' \leftarrow \text{sp}$ 
17:     $\text{sp} \leftarrow 1 + \text{fSum}(\text{lca}(\ell_{\text{sp}}, \ell_{\text{ep}}))$ 
18:     $\text{ep} \leftarrow \text{sp} + \text{rangeCount}(\text{sp}', \text{ep}, d + 1, \sigma_p)$ 
19:  else
20:     $f \leftarrow$  first occurrence of  $c$  in  $P[i, p]$ 
21:     $d \leftarrow$  number of distinct p-characters in  $P[i, f]$ 
22:     $\text{sp}' \leftarrow \text{sp}$ 
23:     $\text{sp} \leftarrow \text{pLF}(\min\{j \mid j \in [\text{sp}, \text{ep}] \text{ and } \text{pBWT}[j] = d\})$ 
24:     $\text{ep} \leftarrow \text{sp} + \text{rangeCount}(\text{sp}', \text{ep}, d, d)$ 
25:  end if
26: end while
27: if ( $\text{sp} < \text{ep}$ ) then “no match found” else return  $[\text{sp}, \text{ep}]$ 

```

using a `rangeNextVal`-query. By comparing β with $\text{pBWT}[i]$, we know whether $u_2 = \text{zeroNode}(\ell_i)$, or not. Note that the α values of all nodes in the tree can be encoded in $\mathcal{O}(n)$ bits using unary encoding, such that each α value can be decoded in $\mathcal{O}(1)$ time. By storing the `zeroDepth` of a carefully selected set of $2^{\lceil \frac{n}{\log \sigma} \rceil}$ marked nodes, we can find a marked node w , which is at most $\log \sigma$ nodes above $\text{zeroNode}(\ell_i)$. Use `rangeNextVal` queries at each node on the path from w to ℓ_i to find $\text{zeroNode}(\ell_i)$. The space-and-time tradeoff is $\mathcal{O}(n)$ bits vs $\mathcal{O}(\log^2 \sigma)$ time; details are presented in Section 6.1. A more sophisticated scheme in Section 6.2 enables us to achieve the claimed bounds. We begin with the following lemma.

Lemma 6 *By maintaining the WT over pBWT and an additional $\mathcal{O}(n)$ -bit structure, in $\mathcal{O}(\log \sigma)$ time, we can find the lowest marked ancestor w_i of a leaf ℓ_i such that $\text{zeroDepth}(w_i) < \text{pBWT}[i]$.*

Proof: Identify marked nodes with $g = \lceil \log \sigma \rceil$. Maintain a bit-array B such that $B[k] = 1$ iff the node with pre-order rank k is a marked node. Also, maintain a `rank-select` structure on B . The space needed is $\mathcal{O}(n)$ bits. We also maintain an array D , such that $D[k]$ equals the `zeroDepth` of the marked node corresponding to the k th 1-bit in B . Given a marked node with pre-order rank k' ,




Figure 5: Illustration of finding `zeroNode` using $\mathcal{O}(n)$ bits

its corresponding position in D is given by $\text{rank}_B(k', 1)$. We do not maintain D explicitly; instead, we maintain a wavelet tree over it. The space needed is $\mathcal{O}(\frac{n}{g} \log \sigma) = \mathcal{O}(n)$ bits.

Given a leaf node ℓ_i , in constant time, first locate the lowest marked ancestor u (refer to Fact 2). Then, find the position j corresponding to u in the array D . If $\text{zeroDepth}(u) < \text{pBWT}[i]$, then $w_i = u$, and we are done. Otherwise, locate $j' = \text{predecessor}_D(j, \text{pBWT}[i])$ i.e., the rightmost position $j' < j$ in D such that $D[j'] < \text{pBWT}[i]$; this is achieved in $\mathcal{O}(\log \sigma)$ time. (Since the root node is marked, and its `zeroDepth` equals 0, the position j' necessarily exists.) Obtain the marked node v corresponding to the j' th 1-bit in B via a $\text{select}_B(j', 1)$ operation. Then, $w_i = \text{lca}(u, v)$. The time required is $\mathcal{O}(\log \sigma)$.

To see the correctness first observe that $\text{lca}(u, v)$ is marked, and is an ancestor of ℓ_i . Also, for any node x , $\text{zeroDepth}(x) \geq \text{zeroDepth}(\text{parent}(x))$. Thus, $\text{zeroDepth}(\text{lca}(u, v)) \leq \text{zeroDepth}(v) < \text{pBWT}[i]$. If $\text{lca}(u, v)$ is not the desired node, then it has a marked descendant $u' \neq u$ on the path to u that satisfies $\text{zeroDepth}(u') < \text{pBWT}[i]$. But then u' appears after v and before u in pre-order, a contradiction. ■

6.1 An $\mathcal{O}(n)$ -bit and $\mathcal{O}(\log^2 \sigma)$ -time Index

Mark nodes w.r.t a parameter $g = \lceil \log \sigma \rceil$. To verify whether a node is marked (w.r.t g), we maintain a bit-vector B where an entry j is set to 1 iff the node with pre-order rank j is marked. (We may simply use the bit-vector B of Lemma 6.) Using Lemma 6, we first identify the lowest marked ancestor w_i of ℓ_i that satisfies $\text{zeroDepth}(w_i) < \text{pBWT}[i]$ in $\mathcal{O}(\log \sigma)$ time. Consider the nodes $w_i = u_1, u_2, \dots, u_k = \ell_i$ on the w_i - ℓ_i path. For any node $u_{k'}$, $1 \leq k' \leq k$, denote by $[L_{k'}, R_{k'}]$ the range of leaves in its subtree. See Figure 5 for an illustration. Let $\alpha_{k'}$, $k' < k$, be the number of leaves $j \in [L_{k'+1}, R_{k'+1}]$ that satisfy $\text{zeroDepth}(u_{k'}) < \text{pBWT}[j] \leq \text{zeroDepth}(u_{k'+1})$. Using the WT over pBWT , we find $\beta_1 =$ the α_1 th smallest value in $\text{pBWT}[L_2, R_2]$ that is greater than $\text{zeroDepth}(u_1)$

i.e., $\beta_1 = \text{rangeNextVal}(L_2, R_2, \text{zeroDepth}(u_1), \alpha_1)$. If $\text{pBWT}[i] \leq \beta_1$, then clearly $u_2 = \text{zeroNode}(\ell_i)$. Otherwise, by the definition of α_1 and β_1 , we must have $\text{pBWT}[i] > \text{zeroDepth}(u_2) \geq \beta_1$. We again query to obtain $\beta_2 = \text{rangeNextVal}(L_3, R_3, \beta_1, \alpha_2)$. Observe that crucially β_2 is the same as $\text{rangeNextVal}(L_3, R_3, \text{zeroDepth}(u_2), \alpha_2)$. We again compare $\text{pBWT}[i]$ and β_2 to determine whether $u_3 = \text{zeroNode}(\ell_i)$, or not; in the latter case, we query again. Repeat this process until we find $\text{zeroNode}(\ell_i)$. By the definition of $w_i = u_1$, any marked node u_r on the path from u_2 to ℓ_i must satisfy $\text{zeroDepth}(u_r) \geq \text{pBWT}[i]$, and we will find $\text{zeroNode}(\ell_i)$ by the time we reach u_r . (Note that $\text{zeroNode}(\ell_i)$ may be the same as u_r .) If u_1 is a lowest marked node, then we will find $\text{zeroNode}(\ell_i)$ by the time we reach ℓ_i . (Note that $\text{zeroNode}(\ell_i)$ may be the same as ℓ_i .) In either case, the number of nodes traversed is at most g (refer to Fact 2). Therefore, the time required is $\mathcal{O}(g \log \sigma) = \mathcal{O}(\log^2 \sigma)$.

In order to find $\alpha_{k'}$, first observe that each leaf ℓ_j contributes to exactly one f_j on a certain edge, which in turn contributes exactly one to the $\alpha_{k'}$ value for this edge. Therefore, crucially, the total of all $\alpha_{k'}$ values over the entire tree is at most n . Using unary encoding, the values over all edges can be stored in $\mathcal{O}(n)$ bits, and the value of an edge can be accessed in constant time¹³. This completes the proof of the $\mathcal{O}(n)$ space and $\mathcal{O}(\log^2 \sigma)$ time solution.

6.2 Improving the time to $\mathcal{O}(\log \sigma \log \log \sigma)$

To improve the query time, we maintain another group of marked nodes w.r.t $g' = \lceil \log \log \sigma \rceil$. Also maintain a bit-vector B' to identify whether a node is marked, or not. Now, given the lowest g' -marked ancestor w'_i of ℓ_i that satisfies $\text{zeroDepth}(w'_i) < \text{pBWT}[i]$, we now need to traverse and query at $\leq g'$ nodes. Thus, the time improves to $\mathcal{O}(\log \sigma \log \log \sigma)$. However, we need an alternative way to find the node w'_i as maintaining their zeroDepth in a WT (as in Lemma 6) will need $\mathcal{O}(\frac{n}{g'} \log \sigma) = \mathcal{O}(n \log \sigma / \log \log \sigma)$ bits, which violates the space bound.

We can find β_1 (corresponding to u_2 as described above) with a single query costing $\mathcal{O}(\log \sigma)$ time. For any k' , $1 < k' < k$, the successive queries starting from u_3 to $u_{k'+1}$ gives us $\beta_{k'}$. Observe that $\beta_{k'}$ is exactly the $\gamma_{k'}$ th smallest number in $\text{pBWT}[L_{k'+1}, R_{k'+1}]$ which is greater than β_1 , where $\gamma_{k'}$ equals the number of leaves $j \in [L_{k'+1}, R_{k'+1}]$ that satisfy $\beta_1 < \text{pBWT}[j] \leq \text{zeroDepth}(u_{k'+1})$. Alternatively, $\gamma_{k'}$ equals the number of leaves $j \in [L_{k'+1}, R_{k'+1}]$ such that f_j lies on the path from u_2 to $u_{k'+1}$. Therefore, we can find $\beta_{k'}$ with a single $\text{rangeNextVal}(L_{k'+1}, R_{k'+1}, \beta_1, \gamma_{k'})$ -query, provided that the value of $\gamma_{k'}$ is known.

Consider the traversal from u_2 to $u_k = \ell_i$. If u_1 is a lowest marked node, then $|\text{leaf}(u_2) \setminus \text{leaf}(u_{k'})| \leq 2g$, $1 < k' \leq k$ (refer to Fact 2). Otherwise, u_2 is a prime node; as before, let u_r be the highest g -marked node on the path below u_1 . We have $\text{zeroDepth}(u_r) \geq \text{pBWT}[i]$, and $|\text{leaf}(u_2) \setminus \text{leaf}(u_{k'})| \leq 2g$, $1 < k' < r$ (refer to Fact 2). If $r \leq g'$, then we can directly use the solution of Section 6.1; query time will be $\mathcal{O}(\log \sigma \log \log \sigma)$ as we will carry out at most g' rangeNextVal -queries. If $r > g'$, then we will encounter at least one g' -marked node in the traversal (refer to Fact 2). Among these g' -marked nodes lies the node w'_i , which we intend to find.

We show the scenario when u_r exists; the case where u_1 is a lowest marked node follows similarly. Let us first concentrate on how to find γ_s for a g' -marked node u_{s+1} on the path from u_2 to u_{r-1} .

¹³ Create a binary string S , initialized to 0, as follows. For each non-root node u in a pre-order traversal of pST , append to S as many 1s as the value of $\alpha_{k'}$ on the edge to its parent, followed by a 0. Maintain a rank-select structure over S . The value of an edge $(\text{parent}(u), u)$, where u has pre-order rank $k > 1$, is the number of 1s between the k th 0 and the $(k-1)$ th 0. The value is found by using two select operations followed by two rank operations.

Consider the leaves in $\mathcal{L}_s = \text{leaf}(u_2) \setminus \text{leaf}(u_{s+1})$. Among these at most $2g$ leaves, we maintain δ_s = the number of leaves $j \in \mathcal{L}_s$ for which f_j lies on the path from u_2 to u_{s+1} . Note that $\delta_s \leq 2g = 2\lceil \log \sigma \rceil$ can be maintained in $\mathcal{O}(\log \log \sigma)$ bits. Thus, the space for maintaining δ_s values over all g' -marked nodes is bounded by $\mathcal{O}(\frac{n}{g} \log \log \sigma) = \mathcal{O}(n)$ bits. (Basically, for each g' -marked node we store the corresponding δ value w.r.t its lowest g -prime ancestor.) By definition, $\gamma_s = (\alpha_2 + \alpha_3 + \dots + \alpha_s) - \delta_s$. The summation part of γ_s is obtained by walking on the u_2 - u_{s+1} path and querying at each edge (which requires $\mathcal{O}(1)$ time per edge using the unary encoding as discussed before). Since at most g nodes are encountered, the time needed is $\mathcal{O}(\log \sigma)$.

We are now equipped to find the desired node w'_i . Traverse the path from u_2 to ℓ_i (using $\text{levelAncestor}(\ell_i, \cdot)$ -queries) and keep track of the g' -marked nodes encountered until we hit the g -marked node u_r (using the bit-vectors B and B'). Now, we can binary search on this g' -marked nodes to find w'_i . Specifically, for any g' -marked u_{s+1} , we first find γ_s and then use it to find β_s ; the time required is $\mathcal{O}(\log \sigma)$. Now, by comparing β_s with $\text{pBWT}[i]$, we know whether we should proceed to the right (i.e., below u_{s+1}) or to the left (i.e., above u_{s+1}) to locate w'_i . The number of searches is $\mathcal{O}(\log g) = \mathcal{O}(\log \log \sigma)$. Therefore, the total time needed is $\mathcal{O}(\log \sigma \log \log \sigma)$, as desired.

7 Structural Suffix Trees

Shibuya considered the following variant of Problem 1, known as *Structural Matching*.

Problem 2 ([Shi00]) Let Σ be an alphabet of size $\sigma \geq 2$, which is the union of two disjoint sets: Σ_s having σ_s static characters (s -characters) and Σ_p having σ_p parameterized characters (p -characters). For each p -character, we associate a p -character (called the complement character). Two strings S and S' , having the same length, are a structural-match (s -match) iff

- S and S' are a p -match, and
- if a p -character x in S is renamed to y in S' , then the complement (if exists) of x in S is renamed to the complement of y in S' .

Let \mathbb{T} be a text having n characters chosen from Σ . We assume that \mathbb{T} terminates in an s -character $\$$ that appears only once. The task is to index \mathbb{T} , such that for a pattern P (also over Σ), we can report all starting positions (occurrences) of the substrings of \mathbb{T} that are an s -match with P .

Consider the following examples. Let $\Sigma_s = \{A, B, C\}$ and $\Sigma_p = \{w, x, y, z\}$, where the complement pairs are x - w and y - z . Then $AxB y Cx$ is an s -match with $AyBxCy$; in this case, there are no complementarity requirements i.e., it is simply a p -match. Also, $AxBwCx$ is an s -match with $AzByCz$; here, x is paired with z , and w (complement of x) is paired with y (complement of z). However, $AxBwCx$ is not an s -match with $AzBxCz$ (even though they are a p -match); this is because as x is paired with z , w should have been paired with y .

Shibuya presented a $\Theta(n \log n)$ -bit and $\mathcal{O}(|P| \log \sigma + \text{occ})$ time index for this problem. We present the following new results.

Theorem 3 All occ positions where P is an s -match with \mathbb{T} can be found as follows: (a) in $\mathcal{O}(|P| \log \sigma + \text{occ} \cdot \log n)$ time using an $\mathcal{O}(n \log \sigma)$ -bit index, and (b) in $\mathcal{O}((|P| + \text{occ}) \log \sigma \log \log \sigma)$ time using an $n \log \sigma + \mathcal{O}(n)$ -bit index.

We encode a string S as $\text{compl}(S)$ as follows. If $S[i]$ is static, then $\text{compl}(S)[i] = S[i]$. Consider a p-character $S[i]$ and let $j^+ < i$ and $j^- < i$ be the rightmost occurrence of $S[i]$ and the complement of $S[i]$ in $S[1, i-1]$. If there is no occurrence j^+ (resp. j^-), we let $j^+ = -1$ (resp. $j^- = -1$). If $j^+ = j^- = -1$, then replace $S[i]$ by 0. Otherwise, if $j^+ > j^-$, then $\text{compl}(S)[i] = (i - j^+)$. Otherwise, if $j^- > j^+$, then $\text{compl}(S)[i] = -(i - j^-)$. For example, $\text{compl}(AxByCx) = A0B0C4$ and $\text{compl}(AxBwAwCx) = A0B(-2)A2C(-2)A2$. The following lemma is trivially true.

Lemma 7 *Two strings S and S' are an s-match iff $\text{compl}(S) = \text{compl}(S')$.*

The structural suffix tree (sST) is the compact trie of all the strings in the set $\{\text{compl}(\mathbf{T}[i, n] \mid 1 \leq i \leq n\}$. We analogously define the *locus* and the *suffix range* of a pattern in this case.

Consider each circular suffix of \mathbf{T} . Sort them based on their $\text{compl}(\cdot)$ encoding, where character precedence is determined by Convention 1. Then, obtain the last character $L[i]$ of the i th lexicographically smallest suffix. Define $\text{sLF}(i) = \text{sSA}^{-1}[\text{sSA}[i] - 1]$, where $\text{sSA}[\cdot]$ and $\text{sSA}^{-1}[\cdot]$ denotes the suffix array and inverse suffix array value according to $\text{compl}(\cdot)$. Denote by f_i^+ (resp. f_i^-) the first occurrence of $L[i]$ (resp. the complement of $L[i]$) in the circular suffix \mathbf{T}_i . In case, there is no occurrence of $L[i]$'s complement, we take $f_i^- = n + 1$. We define the structural BWT as follows.

$$\text{sBWT}[i] = \begin{cases} L[i], & \text{if } L[i] \in \Sigma_s, \\ \text{number of distinct p-characters in } \mathbf{T}_{\text{sSA}[i]}[1, f_i^+], & \text{if } L[i] \in \Sigma_p \text{ and } f_i^+ < f_i^-, \\ -\text{number of distinct p-characters in } \mathbf{T}_{\text{sSA}[i]}[1, f_i^-], & \text{if } L[i] \in \Sigma_p \text{ and } f_i^+ > f_i^-. \end{cases}$$

Observation 2 *For any $1 \leq i \leq n$, $\text{compl}(\mathbf{T}_{\text{pSA}[i]-1}) =$*

$$\begin{cases} \text{sBWT}[i] \circ \text{compl}(\mathbf{T}_{\text{sSA}[i]}[1, n-1]), & \text{if } \text{sBWT}[i] \in \Sigma_s, \\ 0 \circ \text{compl}(\mathbf{T}_{\text{sSA}[i]}[1, f_i^+ - 1] \circ f_i^+ \circ \text{compl}(\mathbf{T}_{\text{sSA}[i]}[f_i^+ + 1, n-1]), & \text{if } \text{sBWT}[i] \in [1, \sigma_p] \\ 0 \circ \text{compl}(\mathbf{T}_{\text{sSA}[i]}[1, f_i^- - 1] \circ -f_i^- \circ \text{compl}(\mathbf{T}_{\text{sSA}[i]}[f_i^- + 1, n-1]), & \text{if } \text{sBWT}[i] \in [-\sigma_p, -1] \end{cases}$$

It follows from Observation 2 that for two suffixes i and j , of which at least one is preceded by an s-character, determining whether $\text{sLF}(i) < \text{sLF}(j)$, or not, remains same as in Lemma 2. So, we concentrate on the scenario where both $L[i]$ and $L[j]$ are p-characters.

Lemma 8 *Assume $i < j$ and both $L[i]$ and $L[j]$ are parameterized. Let u be the lowest common ancestor of ℓ_i and ℓ_j in pST, and z be the number of 0's in the string $\text{path}(u)$. Then,*

(1) *If $|\text{sBWT}[i]|, |\text{sBWT}[j]| \leq z$, then $\text{sLF}(i) < \text{sLF}(j)$ iff one of the following holds:*

- (a) $\text{sBWT}[i], \text{sBWT}[j] > 0$ and $\text{sBWT}[i] \geq \text{sBWT}[j]$
- (b) $\text{sBWT}[i] < 0 < \text{sBWT}[j]$
- (c) $\text{sBWT}[i], \text{sBWT}[j] < 0$ and $|\text{sBWT}[i]| \leq |\text{sBWT}[j]|$

(2) *If $|\text{sBWT}[i]| \leq z < |\text{sBWT}[j]|$, then $\text{sLF}(i) < \text{sLF}(j)$ iff $\text{sBWT}[i] < 0$*

(3) *If $|\text{sBWT}[i]| > z \geq |\text{sBWT}[j]|$, then $\text{sLF}(i) < \text{sLF}(j)$ iff $\text{sBWT}[j] > 0$*

(4) *If $|\text{sBWT}[i]|, |\text{sBWT}[j]| > z$, then $\text{sLF}(i) > \text{sLF}(j)$ iff one of the following holds:*

- (a) $\text{sBWT}[i] = z + 1$, the leading character on the path from u to ℓ_i is 0, and the leading character on the path from u to ℓ_j is not an s -character.
- (b) $\text{sBWT}[j] = -(z + 1)$, and the leading character on the path from u to ℓ_j is 0.

The sBWT can be maintained in $n \log \sigma + \mathcal{O}(n)$ bits to support the same operations outlined in Section 4.1.1 in $\mathcal{O}(\log \sigma)$ time¹⁴.

7.1 Redefining $\text{zeroNode}(\ell_i)$

For a leaf ℓ_i , $\text{sBWT}[i] \in [1, \sigma_p]$, we define $\text{zeroNode}^+(\ell_i)$ to be the locus of $\text{path}(\ell_i)[1, f_i^+]$. Equivalently, $\text{zeroNode}^+(\ell_i)$ is the highest node w on the root to ℓ_i path such that $\text{zeroDepth}(w) \geq \text{sBWT}[i]$. Likewise, for a leaf ℓ_i , $\text{sBWT}[i] < 0$, $\text{zeroNode}^-(\ell_i)$ is the locus of $\text{path}(\ell_i)[1, f_i^-]$. Equivalently, $\text{zeroNode}^-(\ell_i)$ is the highest node w on the root to ℓ_i path such that $\text{zeroDepth}(w) \geq -\text{sBWT}[i]$.

To locate $\text{zeroNode}^+(\ell_i)$ or $\text{zeroNode}^-(\ell_i)$, it suffices to find the lowest node u on the root to ℓ_i path such that $\text{zeroDepth}(u) < |\text{sBWT}[i]|$. (The desired node w is the child of u on this path.)

First, let us concentrate on modifying Lemma 3. We locate the node v with pre-order rank $\text{predecessor}_D(j, |\text{sBWT}[i]|)$ query; here, $D[k]$ equals the zeroDepth of the node with pre-order rank k , and j is the pre-order rank of ℓ_i . Then $u = \text{lca}(\ell_i, v)$. The correctness and complexity claims remain the same as in the case of Lemma 3.

Things are a little bit trickier when it comes to modifying Lemma 5. First of all, we have to find a g -marked node w_i of ℓ_i , where $g = \lceil \log \sigma \rceil$, such that $\text{zeroDepth}(w_i) < |\text{sBWT}[i]|$. This part is achieved by simply replacing the query $\text{predecessor}_D(j, \text{pBWT}[i])$ in the proof of Lemma 6 with the query $\text{predecessor}_D(j, |\text{sBWT}[i]|)$. The correctness and complexity claims remain the same as in the case of Lemma 6. We omit details and definition of the notations in what follows, as they are either the same or follow immediately from the proof of Lemma 5.

As before, $\beta_0 = \text{zeroDepth}(w_i)$. For $k' > 0$, $\beta_{k'}$ is the $\alpha_{k'}$ th smallest $|\text{sBWT}[j]|$ value, which is greater than $\beta_{k'-1}$, $L_{k'+1} \leq j \leq R_{k'+1}$. Here, $\alpha_{k'}$ is the number of leaves j , $L_{k'+1} \leq j \leq R_{k'+1}$, such that either f_j^+ or f_j^- (depending on whether $\text{pBWT}[j] \in [1, \sigma_p]$ or $\text{pBWT}[j] < 0$) lies on the edge from $u_{k'+1}$ to its parent. In the current scenario, starting from u_2 , we have to find the first $\beta_{k'} \geq |\text{sBWT}[i]|$ corresponding to a node $u_{k'+1}$ on the path $w_i = u_1, u_2, \dots, u_k = \ell_i$.

We write $\alpha_{k'}$ separately as $\alpha_{k'}^+$ and $\alpha_{k'}^-$ (corresponding to f_j^+ and f_j^-). All these values can be maintained using unary encoding in $\mathcal{O}(n)$ bits, such that a $\alpha_{k'}^+$ or $\alpha_{k'}^-$ value can be retrieved in $\mathcal{O}(1)$ time. Observe that crucially the desired value $\beta_{k'}$ is the maximum of the $\alpha_{k'}^+$ th smallest value greater than $\beta_{k'-1}$ or the absolute value of the $\alpha_{k'}^-$ th largest value smaller than $-\beta_{k'-1}$. We find $\beta_{k'}^+ = \text{rangeNextVal}(L_{k'+1}, R_{k'+1}, \beta_{k'-1}, \alpha_{k'}^+)$ and $\beta_{k'}^- = \text{rangePrevVal}(L_{k'+1}, R_{k'+1}, -\beta_{k'-1}, \alpha_{k'}^-)$. Then, $\beta_{k'} = \max\{\beta_{k'}^+, -\beta_{k'}^-\}$ is obtained in time $\mathcal{O}(\log \sigma)$.

Based on the above discussion we conclude, $\beta_{k'}^+ = \text{rangeNextVal}(L_{k'+1}, R_{k'+1}, \beta_1, \gamma_{k'}^+)$ and $\beta_{k'}^- = \text{rangePrevVal}(L_{k'+1}, R_{k'+1}, -\beta_1, \gamma_{k'}^-)$, where $\gamma_{k'}^+$ (resp. $\gamma_{k'}^-$) is the number of leaves $j \in [L_{k'+1}, R_{k'+1}]$ such that f_j^+ (resp. f_j^-) lies on the path from u_2 to $u_{k'+1}$. Now, define $\delta_{k'}^+$ as the number of leaves j in $\mathcal{L}_{k'} = \text{leaf}(u_2) \setminus \text{leaf}(u_{k'+1})$ such that f_j^+ lies on the path from u_2 to $u_{k'+1}$. Likewise, define $\delta_{k'}^-$. Both can be maintained in $\mathcal{O}(\log \log \sigma)$ bits at each g' -marked node, where $g' = \lceil \log \log \sigma \rceil$. The total space needed is $\mathcal{O}(n)$ bits. Now, for a g' -marked node on the path from w_i to $u_{k'+1}$,

¹⁴ Create an array sBWT' such that $\text{sBWT}'[i] = \text{sBWT}[i] + \sigma_p$. The alphabet size is at most $\sigma + \sigma_p \leq 2\sigma$. Therefore, the WT over sBWT' occupies space $n \log(2\sigma) + o(n) = n \log \sigma + \mathcal{O}(n)$ bits. A query on sBWT can be easily answered by the same query on sBWT' with σ_p added to the relevant parameters. Time needed is $\mathcal{O}(\log(2\sigma)) = \mathcal{O}(\log \sigma)$.

$\gamma_s^+ = (\alpha_2^+ + \alpha_3^+ + \dots + \alpha_s^+) - \delta_s^+$ is obtained by walking on the path and querying at each node. Likewise, we can find γ_s^- . The time needed is $\mathcal{O}(\log \sigma)$ as we need to walk at most $\lceil \log \sigma \rceil$ nodes. Therefore, as in Lemma 5, we can find the desired g' -marked node via a binary search in time $\mathcal{O}(\log \sigma \log \log \sigma)$.

Summarizing, we have the following lemma.

Lemma 9 *For any leaf ℓ_i , where $\text{sBWT}[i] \in [1, \sigma_p]$, we can compute $\text{zeroNode}^+(\ell_i)$ as follows: (a) in $\mathcal{O}(\log \sigma)$ time using an $\mathcal{O}(n \log \sigma)$ -bit structure, and (b) in $\mathcal{O}(\log \sigma \log \log \sigma)$ time using an $n \log \sigma + \mathcal{O}(n)$ -bit structure. Likewise, for any leaf ℓ_i , where $\text{sBWT}[i] < 0$, we can compute $\text{zeroNode}^-(\ell_i)$ is the same space-and-time tradeoffs.*

7.2 Redefining $\text{fSum}(x)$

Define $\text{fCount}^+(x)$ as the number of leaves ℓ_j under the subtree of x , $\text{sBWT}[j] \in [1, \sigma_p]$, such that $|\text{path}(y)| + 2 \leq f_j \leq |\text{path}(x)| + 1$, where $y = \text{parent}(x)$. Define $\text{fSum}^+(x)$ as the sum of $\text{fCount}^+(y)$ of all nodes y which come before x in pre-order and are not ancestors of x .

Now, we define $\overleftarrow{\text{fSum}}^-(x)$ as the number of leaves ℓ_j , where $\text{sBWT}[j] < 0$, that appear after $\text{rmostLeaf}(x)$ in pre-order and satisfy $f_j^- > 1 + |\text{path}(\text{lca}(x, \ell_j))|$. Thus, $\overleftarrow{\text{fSum}}^-(x)$ is the sum of $\text{fCount}^-(y)$ of all nodes y which come after $\text{rmostLeaf}(x)$ in pre-order (and by definition, are not ancestors of x).

Recall that in the original computation of $\text{fSum}(x)$, we only encode the information for those leaves ℓ_j , where $\text{pBWT}[j] \in [1, \sigma_p]$, i.e., we simply ignore all other leaves. Therefore, while creating the fSum -computation structure, by only considering the leaves ℓ_j , where $\text{sBWT}[j] \in [1, \sigma_p]$, we can obtain a structure for computing $\text{fSum}^+(x)$. Intuitively, $\overleftarrow{\text{fSum}}^-(x)$ is same as $\text{fSum}^+(x)$ with the following modifications:

- We consider the leaves ℓ_j for which $\text{sBWT}[j] < 0$.
- The leaves in sST are ordered in reverse lexicographic order to that of Convention 1.

Therefore, we can compute $\overleftarrow{\text{fSum}}^-(x)$ using a structure similar to that for computing $\text{fSum}^+(x)$.

Summarizing, we have the following lemma.

Lemma 10 *For any node x , we can compute $\text{fSum}^+(x)$ and $\overleftarrow{\text{fSum}}^-(x)$ in $\mathcal{O}(1)$ time by maintaining an $\mathcal{O}(n)$ -bit structure.*

7.3 Implementing sLF Mapping

If $\text{sBWT}[i] > \sigma_p$, then

$$\text{LF}(i) = 1 + \text{rangeCount}(1, n, -\sigma_p, c - 1) + \text{rangeCount}(1, i - 1, c, c)$$

where $c = \text{sBWT}[i]$. We concentrate on the scenario when $\text{sBWT}[i] \in [-\sigma_p, \sigma_p]$.

The following remain unchanged: (a) $\text{leafLeadChar}(j)$ for a leaf ℓ_j , and (b) $\text{pCount}(x)$ for a node x . First find $z = \text{zeroNode}^+(\ell_i)$ or $z = \text{zeroNode}^-(\ell_i)$ depending on whether $\text{pBWT}[i] \in [1, \sigma_p]$ or $\text{pBWT}[i] < 0$. Now, locate the node $v = \text{parent}(z)$. Depending on whether $\text{leafLeadChar}(i) = 0$, or not, we find the desired ranges $\mathcal{S}_1, \mathcal{S}_2, \mathcal{S}_3$, and if required \mathcal{S}_4 and \mathcal{S}_5 . (See Figure 6 for illustration.) Define N_k to be the number of leaves ℓ_j in \mathcal{S}_k such that $\text{sLF}(j) \leq \text{sLF}(i)$. Then, $\text{sLF}(i) = N_1 + N_2 + N_3 + N_4 + N_5$ is computed as follows.




Figure 6: Illustration of various suffix ranges when the suffix $T_{\text{sSA}[i]}$ is preceded by a p-character

7.3.1 Computing N_1

For any $\ell_j \in \mathcal{S}_1$, $\text{sLF}(j) < \text{sLF}(i)$ iff one of the following holds:

- $\text{sBWT}[j] \in [1, \sigma_p]$ and $f_j^+ > 1 + |\text{path}(\text{lca}(z, \ell_j))|$
- $\text{sBWT}[j] < 0$

Therefore,

$$N_1 = \begin{cases} \text{fSum}^+(v) + \text{rangeCount}(1, L_v - 1, -\sigma_p, -1), & \text{if } \text{leafLeadChar}(i) = 0, \\ \text{fSum}^+(z) + \text{rangeCount}(1, L_z - 1, -\sigma_p, -1), & \text{otherwise.} \end{cases}$$

7.3.2 Computing N_2

1. Assume $\text{sBWT}[i] \in [1, \sigma_p]$. For any leaf $\ell_j \in \mathcal{S}_2$, $\text{sLF}(j) \leq \text{sLF}(i)$ iff one of the following holds:

- $\text{sBWT}[j] \in [1, \sigma_p]$ and either $\text{sBWT}[j] > \text{sBWT}[i]$ or $\text{sBWT}[j] = \text{sBWT}[i]$ and $j \leq i$
- $\text{sBWT}[j] < 0$

Therefore,

$$N_2 = \text{rangeCount}(L_z, R_z, c + 1, \sigma_p) + \text{rangeCount}(L_z, i, c, c) + \text{rangeCount}(L_z, R_z, -\sigma_p, -1)$$

where $c = \text{sBWT}[i]$.

2. Assume $\text{sBWT}[i] < 0$. For any leaf $\ell_j \in \mathcal{S}_2$, $\text{sLF}(j) \leq \text{sLF}(i)$ iff one of the following holds:

- $-1 \geq \text{sBWT}[j] > \text{sBWT}[i]$
- $\text{sBWT}[j] = \text{sBWT}[i]$ and $j \leq i$

Therefore, $N_2 = \text{rangeCount}(L_z, R_z, c + 1, -1) + \text{rangeCount}(L_z, i, c, c)$, where $c = \text{sBWT}[i]$.

7.3.3 Computing N_3

For any leaf $\ell_j \in \mathcal{S}_3$, $\text{sLF}(j) < \text{sLF}(i)$ iff $\text{pBWT}[j] < 0$ and $f_j^- \leq 1 + |\text{path}(\text{lca}(z, \ell_j))|$. Therefore,

$$N_3 = \begin{cases} \text{rangeCount}(1, R_v + 1, -\sigma_p, -1) - \overleftarrow{\text{fSum}}^-(v), & \text{if } \text{leafLeadChar}(i) = 0, \\ \text{rangeCount}(1, R_z + 1, -\sigma_p, -1) - \overleftarrow{\text{fSum}}^-(z), & \text{otherwise.} \end{cases}$$

7.3.4 Computing N_4

Let u be the $\text{pCount}(v)$ th child of v . Note that for any leaf $\ell_j \in \mathcal{S}_4$ such that $\text{sBWT}[j] \in [1, \sigma_p]$, $f_j^+ \neq |\text{path}(v)| + 1$; otherwise, the suffix j should not have deviated from i at the node v . Likewise, if $\text{sBWT}[j] < 0$, then $f_j^- \neq |\text{path}(v)| + 1$.

1. Assume $\text{sBWT}[i] > 0$. Then, N_4 is the number of leaves ℓ_j in \mathcal{S}_4 satisfying $j \leq R_u$ and one of the following:

- $\sigma_p \geq \text{sBWT}[j] \geq \text{sBWT}[i]$
- $\text{sBWT}[j] < 0$

Therefore, $N_4 = \text{rangeCount}(R_z + 1, R_u, \text{sBWT}[i], \sigma_p) + \text{rangeCount}(R_z + 1, R_u, -\sigma_p, -1)$.

2. Assume $\text{sBWT}[i] < 0$. Then, N_4 is the number of leaves ℓ_j in \mathcal{S}_4 such that $j \leq R_u$ and $-1 \geq \text{sBWT}[j] > \text{sBWT}[i]$ i.e., $N_4 = \text{rangeCount}(R_z + 1, R_u, \text{sBWT}[i] - 1, -1)$.

7.3.5 Computing N_5

Note that for any leaf $\ell_j \in \mathcal{S}_5$ such that $\text{sBWT}[j] \in [1, \sigma_p]$, $f_j^+ \neq |\text{path}(v)| + 1$; otherwise, the suffix j should not have deviated from i at the node v . Likewise, if $\text{sBWT}[j] < 0$, then $f_j^- \neq |\text{path}(v)| + 1$. Also, note that the leading character of the path from v to ℓ_j is negative.

1. Assume $\text{sBWT}[i] > 0$. Then, N_5 is the number of leaves ℓ_j in \mathcal{S}_5 that satisfies one of the following:

- $\sigma_p \geq \text{sBWT}[j] \geq \text{sBWT}[i]$
- $\text{sBWT}[j] < 0$

Therefore, $N_5 = \text{rangeCount}(L_v, L_z - 1, \text{sBWT}[i], \sigma_p) + \text{rangeCount}(L_v, L_z - 1, -\sigma_p, -1)$.

2. Assume $\text{sBWT}[i] < 0$. Then, N_5 is the number of leaves ℓ_j in \mathcal{S}_5 such that $-1 \geq \text{sBWT}[j] > \text{sBWT}[i]$ i.e., $N_5 = \text{rangeCount}(L_v, L_z - 1, \text{sBWT}[i] + 1, -1)$.

Summarizing, we obtain the following result.

Theorem 4 *We can compute $\text{sLF}(i)$ as follows: (a) in $\mathcal{O}(\log \sigma)$ time using $\mathcal{O}(n \log \sigma)$ bits, and (b) in $\mathcal{O}(\log \sigma \log \log \sigma)$ time using $n \log \sigma + \mathcal{O}(n)$ bits.*

7.4 Backward Search

As in Section 5, assume that we have already found the suffix range $[\text{sp}_1, \text{ep}_1]$ of $\text{compl}(Q)$. We have to find the suffix range of $\text{compl}(cQ)$. If c is a static character, then

$$\begin{aligned}\text{sp}_2 &= 1 + \text{rangeCount}(1, n, -\sigma_p, c - 1) + \text{rangeCount}(1, \text{sp}_1 - 1, c, c) \\ \text{ep}_2 &= \text{rangeCount}(1, n, -\sigma_p, c - 1) + \text{rangeCount}(1, \text{ep}_1, c, c)\end{aligned}$$

Now, we consider the scenario when c is a parameterized character. We omit most of the details as they are similar to those in Section 5. We have the following two cases.

7.4.1 Case 1 (Neither c nor its complement appears in Q)

Let d be the number of distinct p-characters in Q , which can be computed in $\mathcal{O}(1)$ time. Note that $\text{sLF}(i) \in [\text{sp}_2, \text{ep}_2]$ iff $i \in [\text{sp}_1, \text{ep}_1]$, and either (a) $\text{sBWT}[i] \in [1, \sigma_p]$ and $f_i^+ > |Q|$, or (b) $\text{sBWT}[i] < 0$ and $f_i^- > |Q|$. Then,

$$(\text{ep}_2 - \text{sp}_2 + 1) = \text{rangeCount}(\text{sp}_1, \text{ep}_1, d + 1, \sigma_p) + \text{rangeCount}(\text{sp}_1, \text{ep}_1, -\sigma_p, -d - 1)$$

Let $u = \text{lca}(\ell_{\text{sp}_1}, \ell_{\text{ep}_1})$. For any i , $\text{sLF}(i) < \text{sp}_1$ iff one of the following conditions hold:

- $i < \text{sp}_1$, $\text{sBWT}[i] \in [1, \sigma_p]$ and $f_i^+ > 1 + |\text{path}(\text{lca}(u, \ell_i))|$
- $i \in [\text{sp}_1, \text{ep}_1]$, $\text{sBWT}[i] < 0$ and $f_i^- \leq |Q|$
- $i < \text{sp}_1$ and $\text{sBWT}[i] < 0$
- $i > \text{ep}_1$, $\text{sBWT}[i] < 0$ and $f_i^- \leq 1 + |\text{path}(\text{lca}(u, \ell_i))|$

Therefore,

$$\begin{aligned}\text{sLF}(i) &= 1 + \text{fSum}^+(u) + \text{rangeCount}(\text{sp}_1, \text{ep}_1, -d, -1) + \text{rangeCount}(1, \text{sp}_1 - 1, -\sigma_p, -1) \\ &\quad + \text{rangeCount}(\text{ep}_1 + 1, n, -\sigma_p, -1) - \overleftarrow{\text{fSum}}^-(u)\end{aligned}$$

7.4.2 Case 2 (c or its complement appears in Q)

Assume that the number of characters until the first occurrence of c (resp. c 's complement) in Q is f^+ (resp. f^-). If f^+ or f^- does not exist, we take it to be $|Q| + 1$. Let d^+ and d^- be respectively the number of distinct p-characters in $Q[1, f^+]$ and $Q[1, f^-]$ respectively. It follows from the arguments in Section 5.2 that d^+ and d^- can be retrieved in $\mathcal{O}(1)$ time after an initial $\mathcal{O}(|P| \log \sigma)$ time pre-processing.

- **Case when $f^+ < f^-$:** Note that $\text{sLF}(i) \in [\text{sp}_2, \text{ep}_2]$ iff $i \in [\text{sp}_1, \text{ep}_1]$, $\text{sBWT}[i] \in [1, \sigma_p]$ and $f_i = f^+$. Consider any $i, j \in [\text{sp}_1, \text{ep}_1]$ such that $i < j$, both $\text{sLF}(i), \text{sLF}(j) \in [\text{sp}_2, \text{ep}_2]$, and both $\text{sBWT}[i], \text{sBWT}[j] \in [1, \sigma_p]$. Now, $f_i^+ = f_j^+ = f^+$, and $\text{sLF}(i) < \text{sLF}(j)$. Therefore,

$$\begin{aligned}(\text{ep}_2 - \text{sp}_2 + 1) &= \text{rangeCount}(\text{sp}_1, \text{ep}_1, d^+, d^+), \text{ and} \\ \text{sp}_2 &= \text{sLF}(\min\{j \mid j \in [\text{sp}_1, \text{ep}_1] \text{ and } \text{sBWT}[j] = d^+\})\end{aligned}$$

- **Case when $f^+ > f^-$:** Based on the above arguments, we can similarly derive the following.

$$(\text{ep}_2 - \text{sp}_2 + 1) = \text{rangeCount}(\text{sp}_1, \text{ep}_1, -d^-, -d^-), \text{ and} \\ \text{sp}_2 = \text{sLF}(\min\{j \mid j \in [\text{sp}_1, \text{ep}_1] \text{ and } \text{sBWT}[j] = -d^-\})$$

Finally, we remark that we can suitably modify Lemma 1 in a rather straight-forward way, so as to find the $\text{sSA}[j]$ of any j belonging to the suffix range of $\text{compl}(P)$. The time and space requirements remain unchanged. By appropriately choosing Δ , we arrive at Theorem 3.

8 Dictionary Matching

Idury and Schäffer considered the following variant of Problem 1, known as *Parameterized Dictionary Matching*.

Problem 3 ([IS94]) *Let Σ be an alphabet of size $\sigma \geq 2$, which is the union of two disjoint sets: Σ_s having σ_s static characters (s -characters) and Σ_p having σ_p parameterized characters (p -characters). Let \mathcal{D} be a collection of d patterns $\{P_1, P_2, \dots, P_d\}$ of total length n characters that are chosen from Σ . The task is to index \mathcal{D} , such that given a text T (also over Σ), we can report all the positions j in T where at least one pattern $P_i \in \mathcal{D}$ is a p -match with $T[j - |P_i| + 1, j]$.*

Convention 2 *Without loss of generality, assume that no two patterns P_i and P_j exist such that $\text{prev}(P_i) = \text{prev}(P_j)$. Furthermore, assume that $i < j$ iff $\overleftarrow{\text{prev}}(P_i)$ precedes $\overleftarrow{\text{prev}}(P_j)$ in the lexicographic order, where $\overleftarrow{\text{prev}}(P) = \text{prev}(P[p] \circ P[p-1] \circ \dots \circ P[1])$ and $p = |P|$.*

Idury and Schäffer presented an index, based on the Aho-Corasick (AC) automaton [AC75]. The index occupies $\Theta(m \log m)$ bits, where $m \leq n + 1$ is the number of states in the automaton, and can report all *occ* occurrences in time $\mathcal{O}(|T| \log \sigma + \text{occ})$. Typically, an occurrence implies a pair: the position in the text, as well as a pattern in the dictionary that is a p -match with the same-length substring ending at this position. The following theorem summarizes our contribution.

Theorem 5 *All *occ* pairs $\langle j, P_i \rangle$, where a pattern $P_i \in \mathcal{D}$ is a p -match with $T[j - |P_i| + 1, j]$ can be found as follows:*

- in $\mathcal{O}(|T| \log \sigma + \text{occ})$ time using an $\mathcal{O}(m \log \sigma + d \log(n/d))$ -bit index.*
- in $\mathcal{O}(|T| \log \sigma \log \log \sigma + \text{occ})$ time using an $m \log \sigma + \mathcal{O}(m + d \log(n/d))$ -bit index.*

8.1 Idury and Schäffer's Index

Let us first look at the index of Idury and Schäffer [IS94]. We begin by obtaining $\text{prev}(P_i)$ for every P_i in \mathcal{D} , and then create a trie for all the encoded patterns. The number of nodes in the trie is $m \leq n + 1$. For each node u in the trie, denote by $\text{path}(u)$ the string formed by concatenating the edge labels from root to u . Mark a node u in the trie as *final* iff $\text{path}(u) = \text{prev}(P_i)$ for some P_i in \mathcal{D} . Clearly, the number of final nodes is d .

For any prev -encoded string $\text{prev}(S)$ of a string S , and an integer $j \in [1, |S|]$, we obtain a string $\zeta(S, j)$ as follows. Initialize $\zeta(S, j) = \text{prev}(S)[j, |S|]$. For each $j' \in [1, |S| - j + 1]$, assign $\zeta(S, j)[j'] = 0$ iff $\zeta(S, j)[j'] \geq j'$ and $\zeta(S, j)[j'] \notin \Sigma_s$. Conceptually, $\zeta(S, j) = \text{prev}(S[j, |S|])$.

Each node u is associated with 3 links as defined below:

- $\text{next}(u, c) = v$ iff the label on the edge from the node u to v is labeled by the character c . This transition can be carried out in $\mathcal{O}(1)$ time using a perfect hash function.
- $\text{failure}(u) = v$ iff $\text{path}(v) = \zeta(\text{path}(u), j)$, where $j > 1$ is the smallest index for which such a node v exists. If no such j exists, then $\text{failure}(u)$ points to the root node. Conceptually, this represents the smallest shift to be performed in T in case of a mismatch.
- $\text{report}(u) = v$ iff v is a final node and $\text{path}(v) = \zeta(\text{path}(u), j)$, where $j > 1$ is the smallest index for which such a node v exists. If no such j exists, then $\text{report}(u)$ points to the root node. Conceptually, this represents a pattern which has an occurrence ending at the current symbol of the text.

Summarizing, the total space needed by the index is $\Theta(m \log m)$ -bits. Moving forward, we use the terms node and state interchangeably. Likewise, for links and transitions.

To find the desired occurrences, first obtain $T' = \text{prev}(T)$. Now, match T' in the trie as follows. Suppose, we are considering the position j in T' (initially, $j = 1$), and we are at a node u with $|\text{path}(u)| = \delta_u$ i.e., we have matched $T[j, j + \delta_u - 1]$ in the trie. First, repeatedly follow **report**-links starting from u until the root node is reached. Effectively, we report all patterns with a p-match ending at $(j + \delta_u - 1)$. Now, look at the character $c = T'[j + \delta_u]$ to match. If $T'[j + \delta_u] \geq \delta_u$ and $T'[j + \delta_u] \notin \Sigma_s$, then take $c = 0$. If $\text{next}(u, c)$ is defined then follow it to a node v , and try to match an outgoing edge of v with the character $T'[j + \delta_v]$. Otherwise, follow $\text{failure}(u)$ to a node w . Now, we try to match c with an outgoing edge of w ; however, we are now considering the position $j = j + |\text{path}(u)| - |\text{path}(w)|$. Repeat this process until the last character of T' is reached. The total time required is $\mathcal{O}(|T| \log \sigma + \text{occ})^{15}$.

8.2 Representing States

Belazzougui [Bel10] obtained a succinct representation of the AC automaton for the classical dictionary indexing problem [AC75]. Among other techniques, the central idea is to use a Succinctly Indexable Dictionary (SID); see Fact 3 below. For a state u , let $\overleftarrow{\text{path}}(u)$ denote the reverse of $\text{path}(u)$. A state u is labeled by r_u i.e., the lexicographic rank of $\overleftarrow{\text{path}}(u)$ among the set $\{\overleftarrow{\text{path}}(v) \mid v \text{ is a state in the AC automaton}\}$. A transition $\text{next}(u, c)$ is depicted by the bit representation of the string $c \circ r_u$. By maintaining all representations in an SID, the desired node for a transition $\text{next}(u, c)$ can be computed in $\mathcal{O}(1)$ time. The space needed is $m \log \sigma + \mathcal{O}(m)$ bits.

Fact 3 ([RRR02, RRS07]) *A set \mathcal{S} of k integer keys from a universe of size U can be stored in $k(\log(U/k) + \mathcal{O}(1))$ bits of space to support the following two operations in $\mathcal{O}(1)$ time.*

- *return the key of rank i in the natural order of integers*
- *return the rank of key j in the natural order of integers if $j \in \mathcal{S}$. Otherwise, return -1 .*

An important observation related to the index of Idury and Schäffer [IS94], described in the previous section, is that the edges in the trie are labeled from an alphabet of size $\Theta(m)$ in the worse case

¹⁵ Initially T' is obtained in $\mathcal{O}(|T| \log \sigma)$ time. On following a **report** link, either we report an occurrence, or we reach the root. Following this, either we take a **next** transition or we follow a **failure** link; the number of such operations combined is at most $2|T|$. Since each transition takes $\mathcal{O}(1)$ time, the total time is $\mathcal{O}(|T| \log \sigma + \text{occ})$.

scenario. This proves to be the primary bottleneck that prevents us from directly applying the technique of Belazzougui [Bel10] to obtain a succinct representation. More specifically, the SID in this case will need $m \log m + \mathcal{O}(m)$ bits of space, which is not desirable.

To alleviate this problem, we first modify the labeling in the trie of Idury and Schäffer. Moving forward, we denote the trie by \mathcal{T} . Assign a bit to every edge, which is set to 1 iff the labeling on this edge corresponds to a p-character. This can be easily achieved while building \mathcal{T} . Accordingly, we categorize an edge as a *p-edge* or an *s-edge*. Every p-edge is in one of two states: *visited* or *new*. Initially all p-edges are new. Initialize a counter $C = 1$. Traverse from the root node to the leftmost leaf (the ordering of the leaves can be in arbitrary order), and modify the labeling as follows:

- For an s-edge (u, v) , store the value of C at v , and move to the next edge on the path.
- If a p-edge (u, v) is labeled by 0, then label it by the current value of C . Store the value of C at v , and then increment C by 1. Move to the next edge on the path.
- If a p-edge (u, v) is labeled by $c > 0$, then assign it the same label as the $(c + 1)$ th edge on the path from v to root. Store the value of C at v , and move to the next edge on the path.

After the i th leftmost leaf ℓ_i is reached, find $x = \text{lca}(\ell_i, \ell_{i+1})$, and use the stored value of C at x to label the edges from x to ℓ_{i+1} . By pre-processing the trie with the data structure of Section 4.1.2, the entire process can be carried out in $\mathcal{O}(m)$ time.

Observe that each edge is now labeled by a character from the Σ . Let $\overleftarrow{\text{path}}_\Sigma(u)$ (resp. $\text{path}_\Sigma(u)$) denote the concatenation of the new edge labels on the path from u to root (resp. from root to u). With slight abuse of notation, let $\overleftarrow{\text{prev}}(u) = \text{prev}(\overleftarrow{\text{path}}_\Sigma(u))$. Each state u is *conceptually labeled* by the lexicographic rank of $\overleftarrow{\text{prev}}(u)$ in the set $\{\overleftarrow{\text{prev}}(v) \mid v \text{ is a node in the trie}\}$. Thus, each state is labeled by a number in $[1, m]$, where the root is labeled by 1. To distinguish the final states, we *explicitly store* their labels in a SID; space required is $d \log(m/d) + \mathcal{O}(d)$ bits. Using this, given the label of a final state, we can find the corresponding pattern in $\mathcal{O}(1)$ time (refer to Convention 2 and Fact 3). Lastly, we maintain a bit-vector $L[1, m]$ such that $L[j] = 1$ iff the state with label j is a leaf in the trie. Total space occupied is $m + d \log(m/d) + \mathcal{O}(d)$ bits.

8.3 Handling next transitions

We begin with a few notations. For any p-edge $e = (w, x)$, we define $Z(x) =$ the number of 0's in $\overleftarrow{\text{prev}}(w)[1, f_x]$, where f_x is the first occurrence of the p-character labeling e in the string $\overleftarrow{\text{path}}_\Sigma(w)$. If, f_x is not defined, we let $Z(x)$ equal the number of 0's in $\overleftarrow{\text{prev}}(x)$. For any s-edge $e = (w, x)$, we define $Z(x) =$ the label of the edge e . As before, we map the s-characters to the interval $[\sigma_p + 1, \sigma]$, where the i th smallest s-character has value $(\sigma_p + i)$.

Observation 3 For any node x , $Z(x_i) \neq Z(x_j)$ for any two children x_i and x_j of x . Also, $\overleftarrow{\text{prev}}(x_i)$ is lexicographically smaller than $\overleftarrow{\text{prev}}(x_j)$ iff at least one of the following holds:

- $\sigma_p < Z(x_i) < Z(x_j)$
- $Z(x_i) \leq \sigma_p < Z(x_j)$
- $\sigma_p \leq Z(x_j) < Z(x_i)$

Observation 4 Suppose we are at a node u , and have matched $T[j, j + \delta_u - 1]$. We have to select the edge corresponding to $c = T[j + \delta_u]$. If $c \in \Sigma_s$, then clearly we need to select the s -edge (u, v) such that $Z(v) = c$. Now, assume that $c \in \Sigma_p$.

- If $c \notin T[j, j + \delta_u - 1]$, then clearly we need to select the edge $e = (u, v)$ such that $Z(v)$ equals the number of distinct p -characters in $T[j, j + \delta_u]$.
- Otherwise, let the number of distinct p -characters in $T[j', j + \delta_u - 1]$ be z , where j' is the last occurrence of c in $T[j, j + \delta_u - 1]$. Then, we have to select the edge (u, v) such that $Z(v) = z$.

For any two distinct nodes u and v in \mathcal{T} , we denote $u \prec v$ iff $\overleftarrow{\text{prev}}(u)$ is lexicographically smaller than $\overleftarrow{\text{prev}}(v)$. Since $\overleftarrow{\text{prev}}(u) \neq \overleftarrow{\text{prev}}(v)$, the relation $u \prec v$ is well-defined.

We create a compressed $\overleftarrow{\mathcal{T}}$ as follows. Initially $\overleftarrow{\mathcal{T}}$ is empty. For each non-leaf node u in \mathcal{T} , we add the string $\overleftarrow{\text{prev}}(u) \circ \$_{u,i}$ to $\overleftarrow{\mathcal{T}}$ for each child u_i of u . Clearly, each string corresponds to a leaf, say $\ell_{u,i}$, in $\overleftarrow{\mathcal{T}}$. We order the leaves according to the (lexicographic) rank of the string they represent. The (lexicographic) rank of two strings $\overleftarrow{\text{prev}}(u) \circ \$_{u,i}$ and $\overleftarrow{\text{prev}}(u) \circ \$_{u,j}$ (i.e., of two leaves which share the same parent) are determined according $Z(u_i)$ and $Z(u_j)$ as outlined in Observation 3. If two leaves have distinct parents, then their order is defined by the relation \prec on their parents. In both the cases, the rank of the two strings corresponding to any two leaves is well defined.

Note that the number leaves in $\overleftarrow{\mathcal{T}}$ is same as the number of edges in \mathcal{T} . Therefore, it has $(m-1)$ leaves and at most $(m-2)$ internal nodes. Maintain the array Z , corresponding to the leaves in $\overleftarrow{\mathcal{T}}$, as a Wavelet Tree. Also, maintain a succinct representation of $\overleftarrow{\mathcal{T}}$. The total space required to store this information is $m \log \sigma + \mathcal{O}(m)$ bits. The following lemma is a direct consequence of Observation 1, Lemma 2, and Observation 3.

Lemma 11 Consider two nodes u and v (not necessarily distinct) and its respective children u_i and v_j in \mathcal{T} . Let the respective characters (from Σ) on the edges be c_i and c_j .

- If c_i is parameterized and c_j is static, then $\text{next}(u, c_i) \prec \text{next}(v, c_j)$.
- If both c_i and c_j are static, then $\text{next}(u, c_i) \prec \text{next}(v, c_j)$ iff one of the following holds:
 - $Z(u_i) < Z(v_j)$
 - $Z(u_i) = Z(v_j)$ and $u \prec v$. Note that in this case $u \neq v$ and \prec is defined.
- Assume both c_i and c_j are parameterized and either $u \prec v$ or $u = v$. Let $x = \text{lca}(u, v)$, and z be the number of 0's in $\text{prev}(x)$.
 - If $Z(u_i), Z(v_j) \leq z$, then $\text{next}(u, c_i) \prec \text{next}(v, c_j)$ iff $Z(u_i) \geq Z(v_j)$.
 - If $Z(u_i) \leq z < Z(v_j)$, then $\text{next}(v, c_j) \prec \text{next}(u, c_i)$.
 - If $Z(v_j) \leq z < Z(u_i)$, then $\text{next}(u, c_i) \prec \text{next}(v, c_j)$.
 - If $Z(u_i), Z(v_j) > z$, then $\text{next}(v, c_j) \prec \text{next}(u, c_i)$ iff all of the following are true:
 - $Z(u_i) = z + 1$.
 - the leading character on the path from x to $\ell_{u,i}$ is 0.
 - the leading character on the path from x to $\ell_{v,j}$ is not an s -character.

Crucially, Lemma 11 above is exactly the same as Lemma 2 with the following changes:

- pBWT replaced by Z.
- pLF mapping replaced by next transition.
- the pST replaced by the compressed trie $\overleftarrow{\mathcal{T}}$.

Moreover, recall that in the data structure to compute pLF mapping, we only use the property that pST is a compressed trie (i.e., marking scheme, and the structure for computing fSum and zeroNode only use this property.) We simply redefine $\text{zeroDepth}(\ell_{u,i})$ corresponding to a leaf $\ell_{u,i}$ as the highest node v on the root to $\ell_{u,i}$ path that satisfies $\text{zeroDepth}(v) \geq Z(u_i)$.

Summarizing, in order to compute the next transitions, we construct a structure similar to that for computing pLF mapping. Using this, we can compute a $\text{next}(u, c)$ transition in \mathcal{T} in the same space-and-time trade-off, provided that we know the corresponding leaf and Z entry in $\overleftarrow{\mathcal{T}}$. To this end, we maintain a binary string S as follows. Traverse the leaves from left to right in $\overleftarrow{\mathcal{T}}$, and append a 1-bit for each leaf, followed by a 0-bit whenever we switch from a leaf to the next which has a different parent. Append a 0 both at the beginning and at the end, and maintain a rank-select structure. The space needed is $\mathcal{O}(m)$ bits.

Suppose we are at a node u in \mathcal{T} , and have matched $T[j, j + \delta_u - 1]$, where $\delta_u = |\text{path}(u)|$. We have to find the edge corresponding to the character $c = T[j + \delta_u]$ i.e., find $\text{next}(u, c)$. If u is a leaf (which can be verified by the bit-array L), we know there is no next transition, and we will follow $\text{failure}(u)$. Otherwise, let k be the rank of u . Based on Observation 4, if c is an s-character, then we use it directly; otherwise, we find the encoding z corresponding to c in amortized $\mathcal{O}(\log \sigma)$ time per each character in T^{16} . Locate the leaves ℓ_{sp} and ℓ_{ep} in $\overleftarrow{\mathcal{T}}$, where $\text{sp} = \text{rank}(\text{select}_S(k, 0), 1) + 1$, and $\text{ep} = \text{rank}(\text{select}_S(k + 1, 0), 1)$. The desired leaf and Z entry are obtained as follows. Obtain $q = \text{select}_Z(1 + \text{rank}_Z(\text{sp} - 1, z), z)$ in $\mathcal{O}(\log \sigma)$ time using the WT on Z. If $q > \text{ep}$, then no such leaf exists, and we take the failure link. Otherwise, the desired leaf is given by $\text{child}(\text{parent}(\ell_{\text{sp}}), q - \text{sp} + 1)$. The desired entry in Z is $Z[q]$. To see correctness, observe that ℓ_{sp} and ℓ_{ep} corresponds to the two children u_i and u_j of u such that $\overleftarrow{\text{prev}}(u_i)$ and $\overleftarrow{\text{prev}}(u_j)$ are respectively the lexicographically smallest and largest strings in the set $\{\overleftarrow{\text{prev}}(w) \mid w \text{ is a child of } u\}$. The correctness follows from Observation 3, and the way the leaves under a single node in $\overleftarrow{\mathcal{T}}$ are arranged.

Summarizing, we obtain the following lemma.

Lemma 12 *The next transition in the index of Idury and Schaffer [IS94] can be simulated as follows: (a) in $\mathcal{O}(\log \sigma)$ amortized time per character by using $\mathcal{O}(m \log \sigma)$ bits, and (b) in $\mathcal{O}(\log \sigma \log \log \sigma)$ amortized time per character by using $m \log \sigma + \mathcal{O}(m)$ bits.*

¹⁶ We maintain an array A of length σ_p such that $A[c]$ equals the position of last occurrence of $c \in \Sigma_p$ in the current window $T[j, j + \delta_u - 1]$ of the text. Initially, for each $c \in \Sigma_p$, we assign $A[c] = -1$. Maintain a balanced binary search tree (BST) \mathcal{T}_{bin} containing only the p-characters c indexed by $A[c] \neq -1$. Note that the size of \mathcal{T}_{bin} is $\mathcal{O}(\sigma_p)$, which implies every deletion, insertion, and search operation requires $\mathcal{O}(\log \sigma_p)$ time.

Suppose we are looking at a p-character at position k in the text. If $A[T[k]] = -1$, we add $T[k]$ to \mathcal{T}_{bin} , and find the number of characters in \mathcal{T}_{bin} , which gives us the desired encoding. Otherwise, find the number of entries in \mathcal{T}_{bin} with key at least $A[T[k]]$, which gives us the desired encoding. Finally, update $A[T[k]] = k$ and proceed.

When we follow a failure link as we may have to remove several characters from \mathcal{T}_{bin} . Suppose, after following a failure link, we are considering the window of the text starting from $j' > j$. Clearly, we have to remove those characters c from \mathcal{T}_{bin} for which $A[c] < j'$. Set $A[c] = -1$ for these characters. The total number of such deletion operations is at most $|T|$, yielding an amortized time complexity of $\mathcal{O}(\log \sigma_p)$ per character.

8.4 Handling failure and report transitions

Note that for any two nodes u and v , if $\text{failure}(u) = v$, then it $\overleftarrow{\text{prev}}(v)$ is the longest prefix of $\overleftarrow{\text{prev}}(u)$ that appears in \mathcal{T} . Similar remarks hold for $\text{report}(u) = v$, where v is a final node. Therefore, these behave exactly in the same manner as in the case of traditional pattern matching, and we can simply use the idea of Belazzougui to perform these transitions (see Sections 3.3 and 3.4 in [Bel10]).

Lemma 13 *The failure and report links in the index of Idury and Schäffer [IS94] can be simulated in $\mathcal{O}(1)$ time by using $\mathcal{O}(m + d \log(n/d))$ bits.*

Putting Lemmas 12 and 13 together, we obtain Theorem 5.

9 Discussion

We leave two main questions unanswered, both related to the central problem of parameterized pattern matching. The first one, concerning the space consumption, is “Can we convert $\mathcal{O}(n)$ bits term to $o(n)$ in our space requirements?”. The second one is related to construction of the index. In the case of classical full text indexes, designing linear time construction algorithm in compact space (i.e., the working space is $\mathcal{O}(n \log \sigma)$ bits) has been an active area as seen from the seminal work of Hon, Sadakane, and Sung [HSS03] and the more recent work of Belazzougui [Bel14]. It would be interesting to check for the existence of a (randomized) algorithm for constructing a compressed index for the parameterized matching problem that uses $\mathcal{O}(n \log \sigma)$ working space and attains the same bounds of the best-known algorithms for constructing p-suffix trees [CH03, Kos95].

References

- [AAC⁺03] Amihood Amir, Yonatan Aumann, Richard Cole, Moshe Lewenstein, and Ely Porat. Function matching: Algorithms, applications, and a lower bound. In *Automata, Languages and Programming, 30th International Colloquium, ICALP 2003, Eindhoven, The Netherlands, June 30 - July 4, 2003. Proceedings*, pages 929–942, 2003.
- [AC75] Alfred V. Aho and Margaret J. Corasick. Efficient string matching: An aid to bibliographic search. *Commun. ACM*, 18(6):333–340, 1975.
- [AFM94] Amihood Amir, Martin Farach, and S. Muthukrishnan. Alphabet dependence in parameterized matching. *Inf. Process. Lett.*, 49(3):111–115, 1994.
- [AN09] Amihood Amir and Gonzalo Navarro. Parameterized matching on non-linear structures. *Inf. Process. Lett.*, 109(15):864–867, 2009.
- [Bak93] Brenda S. Baker. A theory of parameterized pattern matching: algorithms and applications. In *Proceedings of the Twenty-Fifth Annual ACM Symposium on Theory of Computing, May 16-18, 1993, San Diego, CA, USA*, pages 71–80, 1993.
- [Bak95] Brenda S. Baker. On finding duplication and near-duplication in large software systems. In *2nd Working Conference on Reverse Engineering, WCRE '95, Toronto, Canada, July 14-16, 1995*, pages 86–95, 1995.
- [Bel10] Djamal Belazzougui. Succinct dictionary matching with no slowdown. In *Combinatorial Pattern Matching, 21st Annual Symposium, CPM 2010, New York, NY, USA, June 21-23, 2010. Proceedings*, pages 88–100, 2010.
- [Bel14] Djamal Belazzougui. Linear time construction of compressed text indices in compact space. In *Symposium on Theory of Computing, STOC 2014, New York, NY, USA, May 31 - June 03, 2014*, pages 148–193, 2014.
- [BM77] Robert S. Boyer and J. Strother Moore. A fast string searching algorithm. *Commun. ACM*, 20(10):762–772, 1977.
- [BM98] Brenda S Baker and Udi Manber. Deducing similarities in java sources from byte-codes. In *USENIX Annual Technical Conference*, pages 179–190, 1998.
- [BW94] M. Burrows and D. J. Wheeler. A block-sorting lossless data compression algorithm. Technical report, 1994.
- [BYdM⁺98] Ira D. Baxter, Andrew Yahin, Leonardo Mendonça de Moura, Marcelo Sant’Anna, and Lorraine Bier. Clone detection using abstract syntax trees. In *1998 International Conference on Software Maintenance, ICSM 1998, Bethesda, Maryland, USA, November 16-19, 1998*, pages 368–377, 1998.
- [CH03] Richard Cole and Ramesh Hariharan. Faster suffix tree construction with missing suffix links. *SIAM J. Comput.*, 33(1):26–42, 2003.

- [CLI] CLoning Analysis and Categorization System (CLICS). <http://www.swag.uwaterloo.ca/clics/>. Accessed: 2015-11-01.
- [DLDPFG01] Giuseppe Antonio Di Lucca, Massimiliano Di Penta, Anna Rita Fasolino, and Pasquale Granato. Clone analysis in the web era: An approach to identify cloned web pages. In *Seventh IEEE Workshop on Empirical Studies of Software Maintenance*, pages 107–113, 2001.
- [dMRS05] Cédric du Mouza, Philippe Rigaux, and Michel Scholl. Efficient evaluation of parameterized pattern queries. In *Proceedings of the 2005 ACM CIKM International Conference on Information and Knowledge Management, Bremen, Germany, October 31 - November 5, 2005*, pages 728–735, 2005.
- [dMRS07] Cédric du Mouza, Philippe Rigaux, and Michel Scholl. Parameterized pattern queries. *Data Knowl. Eng.*, 63(2):433–456, 2007.
- [FH07] Johannes Fischer and Volker Heun. A new succinct representation of rmq-information and improvements in the enhanced suffix array. In *Combinatorics, Algorithms, Probabilistic and Experimental Methodologies, First International Symposium, ESCAPE 2007, Hangzhou, China, April 7-9, 2007, Revised Selected Papers*, pages 459–470, 2007.
- [Fis10] Johannes Fischer. Wee LCP. *Inf. Process. Lett.*, 110(8-9):317–320, 2010.
- [FM00] Paolo Ferragina and Giovanni Manzini. Opportunistic data structures with applications. In *41st Annual Symposium on Foundations of Computer Science, FOCS 2000, 12-14 November 2000, Redondo Beach, California, USA*, pages 390–398, 2000.
- [GGG⁺07] Alexander Golynski, Roberto Grossi, Ankur Gupta, Rajeev Raman, and S. Srinivasa Rao. On the size of succinct indices. In *Algorithms - ESA 2007, 15th Annual European Symposium, Eilat, Israel, October 8-10, 2007, Proceedings*, pages 371–382, 2007.
- [GGV03] Roberto Grossi, Ankur Gupta, and Jeffrey Scott Vitter. High-order entropy-compressed text indexes. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms, January 12-14, 2003, Baltimore, Maryland, USA.*, pages 841–850, 2003.
- [Gia93] Raffaele Giancarlo. The suffix of a square matrix, with applications. In *Proceedings of the Fourth Annual ACM/SIGACT-SIAM Symposium on Discrete Algorithms, 25-27 January 1993, Austin, Texas.*, pages 402–411, 1993.
- [GNP12] Travis Gagie, Gonzalo Navarro, and Simon J. Puglisi. New algorithms on wavelet trees and applications to information retrieval. *Theor. Comput. Sci.*, 426:25–41, 2012.
- [Gus97] Dan Gusfield. *Algorithms on Strings, Trees, and Sequences - Computer Science and Computational Biology*. Cambridge University Press, 1997.

- [GV00] Roberto Grossi and Jeffrey Scott Vitter. Compressed suffix arrays and suffix trees with applications to text indexing and string matching (extended abstract). In *Proceedings of the Thirty-Second Annual ACM Symposium on Theory of Computing, May 21-23, 2000, Portland, OR, USA*, pages 397–406, 2000.
- [HJHC10] B. Hummel, E. Juergens, L. Heinemann, and M. Conradt. Index-based code clone detection: incremental, distributed, scalable. In *Software Maintenance (ICSM), 2010 IEEE International Conference on*, pages 1–9, Sept 2010.
- [HLS04] Carmit Hazay, Moshe Lewenstein, and Dina Sokol. Approximate parameterized matching. In *Algorithms - ESA 2004, 12th Annual European Symposium, Bergen, Norway, September 14-17, 2004, Proceedings*, pages 414–425, 2004.
- [HPST13] Wing-Kai Hon, Manish Patil, Rahul Shah, and Sharma V. Thankachan. Compressed property suffix trees. *Inf. Comput.*, 232:10–18, 2013.
- [HSS03] Wing-Kai Hon, Kunihiko Sadakane, and Wing-Kin Sung. Breaking a time-and-space barrier in constructing full-text indices. In *44th Symposium on Foundations of Computer Science (FOCS 2003), 11-14 October 2003, Cambridge, MA, USA, Proceedings*, pages 251–260, 2003.
- [HSTV14] Wing-Kai Hon, Rahul Shah, Sharma V. Thankachan, and Jeffrey Scott Vitter. Space-efficient frameworks for top- k string retrieval. *J. ACM*, 61(2):9:1–9:36, 2014.
- [IS94] Ramana M. Idury and Alejandro A. Schäffer. Multiple matching of parameterized patterns. In *Combinatorial Pattern Matching, 5th Annual Symposium, CPM 94, Asilomar, California, USA, June 5-8, 1994, Proceedings*, pages 226–239, 1994.
- [JPS13] Markus Jalsenius, Benny Porat, and Benjamin Sach. Parameterized matching in the streaming model. In *30th International Symposium on Theoretical Aspects of Computer Science, STACS 2013, February 27 - March 2, 2013, Kiel, Germany*, pages 400–411, 2013.
- [KEF⁺14] Jinil Kim, Peter Eades, Rudolf Fleischer, Seok-Hee Hong, Costas S. Iliopoulos, Kunsoo Park, Simon J. Puglisi, and Takeshi Tokuyama. Order-preserving matching. *Theor. Comput. Sci.*, 525:68–79, 2014.
- [KFF06] Rainer Koschke, Raimar Falke, and Pierre Frenzel. Clone detection using abstract syntax suffix trees. In *13th Working Conference on Reverse Engineering (WCRE 2006), 23-27 October 2006, Benevento, Italy*, pages 253–262, 2006.
- [KJP77] Donald E. Knuth, James H. Morris Jr., and Vaughan R. Pratt. Fast pattern matching in strings. *SIAM J. Comput.*, 6(2):323–350, 1977.
- [KKL09] Orgad Keller, Tsvi Kopelowitz, and Moshe Lewenstein. On the longest common parameterized subsequence. *Theor. Comput. Sci.*, 410(51):5347–5353, 2009.
- [KKP03] Dong Kyue Kim, Yoo Ah Kim, and Kunsoo Park. Generalizations of suffix arrays to multi-dimensional matrices. *Theor. Comput. Sci.*, 302(1-3):223–238, 2003.

- [Kos95] S. Rao Kosaraju. Faster algorithms for the construction of parameterized suffix trees (preliminary version). In *36th Annual Symposium on Foundations of Computer Science, Milwaukee, Wisconsin, 23-25 October 1995*, pages 631–637, 1995.
- [KR87] Richard M. Karp and Michael O. Rabin. Efficient randomized pattern-matching algorithms. *IBM Journal of Research and Development*, 31(2):249–260, 1987.
- [Lew15] Moshe Lewenstein. Parameterized pattern matching. In *Encyclopedia of Algorithms*. 2015.
- [MM93] Udi Manber and Eugene W. Myers. Suffix arrays: A new method for on-line string searches. volume 22, pages 935–948, 1993.
- [MP15] Juan Mendivelso and Yoan J. Pinzón. Parameterized matching: Solutions and extensions. In *Proceedings of the Prague Stringology Conference 2015, Prague, Czech Republic, August 24-26, 2015*, pages 118–131, 2015.
- [Nav13] Gonzalo Navarro. Spaces, trees, and colors: The algorithmic landscape of document retrieval on sequences. *ACM Comput. Surv.*, 46(4):52:1–52:47, 2013.
- [Nav14] Gonzalo Navarro. Wavelet trees for all. *J. Discrete Algorithms*, 25:2–20, 2014.
- [NM07] Gonzalo Navarro and Veli Mäkinen. Compressed full-text indexes. *ACM Comput. Surv.*, 39(1), 2007.
- [NS14] Gonzalo Navarro and Kunihiro Sadakane. Fully functional static and dynamic succinct trees. *ACM Trans. Algorithms*, 10(3):16:1–16:39, 2014.
- [OGK10] Enno Ohlebusch, Simon Gog, and Adrian Kügel. Computing matching statistics and maximal exact matches on compressed full-text indexes. In *String Processing and Information Retrieval - 17th International Symposium, SPIRE 2010, Los Cabos, Mexico, October 11-13, 2010. Proceedings*, pages 347–358, 2010.
- [RD04] Filip Van Rysselberghe and Serge Demeyer. Evaluating clone detection techniques from a refactoring perspective. In *19th IEEE International Conference on Automated Software Engineering (ASE 2004), 20-25 September 2004, Linz, Austria*, pages 336–339, 2004.
- [RNO11] Luís M. S. Russo, Gonzalo Navarro, and Arlindo L. Oliveira. Fully compressed suffix trees. *ACM Transactions on Algorithms*, 7(4):53, 2011.
- [RRR02] Rajeev Raman, Venkatesh Raman, and S. Srinivasa Rao. Succinct indexable dictionaries with applications to encoding k -ary trees and multisets. In *Proceedings of the Thirteenth Annual ACM-SIAM Symposium on Discrete Algorithms, January 6-8, 2002, San Francisco, CA, USA.*, pages 233–242, 2002.
- [RRS07] Rajeev Raman, Venkatesh Raman, and Srinivasa Rao Satti. Succinct indexable dictionaries with applications to encoding k -ary trees, prefix sums and multisets. *ACM Transactions on Algorithms*, 3(4), 2007.

- [Sad02a] Kunihiro Sadakane. Space-efficient data structures for flexible text retrieval systems. In *Algorithms and Computation, 13th International Symposium, ISAAC 2002 Vancouver, BC, Canada, November 21-23, 2002, Proceedings*, pages 14–24, 2002.
- [Sad02b] Kunihiro Sadakane. Succinct representations of lcp information and improvements in the compressed suffix arrays. In *Proceedings of the Thirteenth Annual ACM-SIAM Symposium on Discrete Algorithms, January 6-8, 2002, San Francisco, CA, USA.*, pages 225–232, 2002.
- [Sad07] Kunihiro Sadakane. Compressed suffix trees with full functionality. *Theory Comput. Syst.*, 41(4):589–607, 2007.
- [Shi00] Tetsuo Shibuya. Generalization of a suffix tree for RNA structural pattern matching. In *Algorithm Theory - SWAT 2000, 7th Scandinavian Workshop on Algorithm Theory, Bergen, Norway, July 5-7, 2000, Proceedings*, pages 393–406, 2000.
- [SN10] Kunihiro Sadakane and Gonzalo Navarro. Fully-functional succinct trees. In *Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2010, Austin, Texas, USA, January 17-19, 2010*, pages 134–149, 2010.
- [Tsu13] Dekel Tsur. Top-k document retrieval in optimal space. *Inf. Process. Lett.*, 113(12):440–443, 2013.
- [Ukk95] Esko Ukkonen. On-line construction of suffix trees. *Algorithmica*, 14(3):249–260, 1995.
- [Wei73] Peter Weiner. Linear pattern matching algorithms. In *14th Annual Symposium on Switching and Automata Theory, Iowa City, Iowa, USA, October 15-17, 1973*, pages 1–11, 1973.